I Still Don't Want to Lose This

Seven months after I Do Not Want To Lose This, a personal return to rented AI capacity, portable thinking tools and who gets to imagine and test alternative futures.

2026.11.03 22:14 DENNIS HEDEGREEN OPEN V1.0
HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/I-STILL-DONT-WANT-TO-LOSE-THIS/

In *I Do Not Want To Lose This*, I wrote about a new kind of access: the ability to take an idea, work on it with a machine, and turn it into something another person could inspect. For someone without an institution behind him, that difference was enormous. [1 (https://hedegreenresearch.com/articles/i-do-not-want-to-lose-this/)]

I still do not want to lose it.

But I need to correct the direction of the story before fear becomes the whole story.

The machine is not finished. Neither is the future it may help us build.

An email about a limit

At the end of October, OpenAI sent good news: the Work and Codex allowance included with Pro 100 had doubled from 29 October 2026. It also described a coming change. Starting in December, ChatGPT and Codex are expected to share a token-based allowance, with more demanding models and tasks consuming more of it. Instant chat is excluded from that shared pool. [2 (https://help.openai.com/en/articles/6825453-chatgpt-release-notes)], [3 (https://help.openai.com/en/articles/20001625-understanding-chatgpt-and-codex-usage-limits)]

That is not evidence that my effective capacity has already been reduced. Part of the current allowance has increased. The effect of the December system on my actual work is still unknown.

And yet the message made something visible.

I have been building a research practice on access that I rent. The provider can change the terms. I can read the announcement, but I cannot calculate its impact on my next paper from the word *tokens* alone.

A token is not a corrected claim. It is not a tested software change. It is not a completed source audit. It is certainly not a measure of public knowledge.

I intend to measure what changes rather than write the conclusion in advance.

But I also want to remember why I cared enough to worry.

What the subscription made possible

Hedegreen Research now contains articles, working papers, methodology, source trails, tools, public records, and a long list of experiments that are not yet ready. There is music too. I have used AI not only to write about the world, but to hear things, build things, and find out whether an idea can survive its first encounter with reality. [4 (https://hedegreenresearch.com/articles/)], [5 (https://hedegreenresearch.com/articles/who-lives-in-signal/)]

That is production. It is not automatically progress.

A paper with a DOI is not automatically good science. An article with references can still misuse them. A program that runs may still be wrong. A song can exist without anybody else valuing it. I have written about the verification debt created when AI makes it too easy to produce convincing work. [6 (https://hedegreenresearch.com/articles/what-if-the-slop-is-mine/)]

In the follow-up to *Nine Became Eight*, I used AI assistance to retrieve later source records, compare nine proof files and their accompanying explanations with the July records, and inspect receipts for another researcher's formal checks. The result became *The World Kept Checking*, with a dated evidence record. [7 (https://hedegreenresearch.com/articles/the-world-kept-checking/)]

Matching files was not the same as proving that the mathematics was new. I had not rerun the proof toolchain, and the full quarterly comparison remained incomplete. The useful output was a narrower, inspectable account of what had been checked and what had not.

I was not given a university. I was not given a newsroom, a development department, or a recording studio.

I was given access to a new kind of production capacity, and I began trying to build some of those functions for myself.

There is an economic question here: has the work created more value than the AI cost?

It is a legitimate question. It also has at least three answers waiting to be confused with one another.

What I paid is one number. What the provider spent serving my work is another. What independent researchers and the public may eventually gain from the outputs is a third. The second number is not public for my account. The third has not been independently measured. Even a deliberately low estimate of replacement labour would not settle it.

I will not claim a proven social return because a spreadsheet produced a pleasing ratio.

But I will not pretend that there is no value here simply because it did not arrive through a conventional institution.

Perhaps PC is the wrong name now

The personal computer made it possible to own computing capacity. AI adds another kind of interaction to that familiar machine.

Now I can discuss an idea with a machine that can propose objections, connect disciplines, write code, interpret sources, and return with a different route than the one I suggested.

I wrote about *The Thinking Machine* in March. Here I want to use the name for a practical category. [8 (https://hedegreenresearch.com/articles/the-thinking-machine/)]

TM. Thinking Machine.

Not because I have established that a model is conscious. I have not. Not because it is always right. It is not. I mean *thinking* operationally: a system able to perform parts of the reasoning, language, planning, and problem-solving work that once required another human participant.

It is no longer only the tool that follows a list of instructions. It can help examine whether the instructions make sense.

That does not make it responsible for my decisions. A machine can help argue; it cannot be used as a place for responsibility to disappear. [9 (https://hedegreenresearch.com/articles/ai-should-be-infrastructure-not-authority/)]

But it changes what one person can attempt.

The door may open from the other side

Today much of the strongest AI capacity is rented through cloud services. That is not the only architecture being built.

NVIDIA already describes local AI development across GeForce RTX, RTX Spark, DGX Spark and larger systems. The DGX Spark class offers up to 128 GB of unified memory; the exact models and workloads it can run still depend on software, precision, speed and hardware configuration. [10 (https://developer.nvidia.com/topics/ai/local-ai)]

The Codex coding agent has a public codebase under the Apache 2.0 licence. That does not make OpenAI's proprietary frontier models open, and it does not mean the best local model can perform every task that a remote model can. [11 (https://github.com/openai/codex/blob/main/LICENSE)]

It does mean the stack is becoming more separable: the software that organizes work, the model that performs it, and the machine that runs it need not always come from the same place.

There is no promised date when an affordable local TM will equal the strongest cloud system. Hardware, electricity, maintenance, model quality and security all remain real costs.

But it is no longer unreasonable to imagine moving more of an independent research operation onto equipment the researcher controls.

That matters more to me than winning an argument about one subscription tier.

If a door narrows, I want to know whether there is another door being built.

Who gets the time?

Owning more of the equipment could change who controls my access. It would not decide who receives the benefit of the work the machine helps perform.

That is the question behind Saved Hours. Time released from a task might become shorter working hours, better care, lower prices or more capacity to investigate. It might also become higher targets and a denser working day. The technical saving does not make that choice. [12 (https://hedegreenresearch.com/articles/the-saved-hours-doctrine/)]

My own output is not a measurement of hours returned: checking, correcting and maintaining the work take time too. Doing more is one possible gain. Having more choice over what to do is another.

This is where Saved Hours meets Human Baseline. If machines can help meet ordinary needs with less required labour, I want to examine whether people can gain more control over their lives. That possibility needs a design and a test, not just a faster machine.

We imagine better machines more easily than better societies

Here is where my optimism becomes less about subscriptions and more about what the machine is for.

We are remarkably comfortable imagining robots that build houses, automated farms, intelligent power grids, and machines that can do more and more of the work people currently perform.

Then, when we discuss the society around them, we often preserve the familiar arrangement and insert a robot into it.

Who employs the robot? Who sells the output? Who pays the rent? Who receives an income? Who owns the bottleneck?

These are necessary questions. They are not the only possible questions.

What if we designed the human requirement first, then tested different ways to meet it?

What would it physically take to keep everyone above a basic food floor? How much energy, land, transport, storage, redundancy, repair and human oversight? Where would the plan fail? Which parts could be cooperative, public, privately supplied or locally operated? [13 (https://hedegreenresearch.com/articles/why-should-anyone-go-hungry/)]

These are not answers that an AI can simply pronounce correct. They are models we can begin to construct, criticize and compare.

That is what interests me about Human Baseline. It is not a proposal to make everything free. Nothing material is without cost. It is an attempt to ask which costs are physically necessary, which are organizational, and which are gates we have chosen to put in front of ordinary life. [14 (https://hedegreenresearch.com/articles/money-should-not-be-the-door-to-survival/)]

I do not want to prohibit people from earning money or collecting expensive things. A first-edition Charizard card can remain rare in a world that has become very good at producing food. Someone can pay a ridiculous amount for it. Someone else can build an extravagant brand. They can trade, speculate, compete and become wealthy.

The question is whether the rest of us must keep buying our survival through the same arrangements, even when alternatives could be built.

That is not the abolition of markets. It is a distinction between the market for preference and the conditions of existence.

It is also not proof that an alternative will be cheaper or fairer. A baseline can be badly governed, expensive, inflexible or captured by a new set of owners. Its advantages have to be demonstrated, not assumed.

The point is that we are allowed to investigate the alternatives.

Intelligence does not decide the destination

A much more capable future AI may be able to model energy systems, production constraints and institutional trade-offs better than we can. It may identify contradictions in the stories societies tell about what is necessary.

It may reach recommendations that are inconvenient for the people financing it. Or it may be used to preserve their advantages.

Neither outcome follows automatically from intelligence.

I do not need a machine to arrive with its own unquestionable morality. I need the machine to help make assumptions visible and consequences testable, including the assumptions I like.

Who defined the objective? Whose freedom was counted? Who pays? Who can say no? Who can repair the outcome when the model was wrong?

Those remain human and institutional questions, even when the reasoning tools become extraordinary.

The work that remains

The March essay was written from the fear that a remarkable door might close. I still need reliable access, predictable costs, and work that survives a change of provider.

But preserving access is only part of the ambition. I want to use it to investigate arrangements that would otherwise remain vague hopes: specify their material requirements, compare who benefits, and find where they break.

Some proposals will fail on water, maintenance, logistics, power, dignity or freedom. We should preserve those failures so the next person does not have to repeat them. The same tools that make an attractive future easier to describe must help make it easier to reject.

That is the opening I care about: more people being able to participate in imagining the future with enough detail to test it.

The first article ended with the need to build things that remain when the economics change. I still mean that. But I would add something now:

Before we can build a different future, we have to be able to imagine one.

And for the first time in my life, I have a machine that can help me do both.

— Dennis Hedegreen, still checking

Sources

[1] Hedegreen Research — I Do Not Want To Lose This (https://hedegreenresearch.com/articles/i-do-not-want-to-lose-this/). 22 March 2026.

[2] OpenAI — ChatGPT release notes (https://help.openai.com/en/articles/6825453-chatgpt-release-notes). 29 October 2026 announcement.

[3] OpenAI — Understanding ChatGPT and Codex usage limits (https://help.openai.com/en/articles/20001625-understanding-chatgpt-and-codex-usage-limits). Announced December transition.

[4] Hedegreen Research — Archive (https://hedegreenresearch.com/articles/). Public work and article records.

[5] Hedegreen Research — Who Lives in Signal (https://hedegreenresearch.com/articles/who-lives-in-signal/). Creative programme and music context.

[6] Hedegreen Research — What If the Slop Is Mine? (https://hedegreenresearch.com/articles/what-if-the-slop-is-mine/). Verification debt and limits of output counts.

[7] Hedegreen Research — The World Kept Checking (https://hedegreenresearch.com/articles/the-world-kept-checking/). Documented AI-mathematics follow-up and dated evidence record; no local proof rebuild or full quarterly comparison.

[8] Hedegreen Research — The Thinking Machine (https://hedegreenresearch.com/articles/the-thinking-machine/). Earlier March reflection.

[9] Hedegreen Research — AI Should Be Infrastructure, Not Authority (https://hedegreenresearch.com/articles/ai-should-be-infrastructure-not-authority/). Human responsibility and institutional control.

[10] NVIDIA Developer — Build Local AI With NVIDIA GPUs (https://developer.nvidia.com/topics/ai/local-ai). Manufacturer descriptions and hardware limits.

[11] OpenAI — Codex repository licence (https://github.com/openai/codex/blob/main/LICENSE). Apache 2.0 software licence; proprietary model weights are separate.

[12] Hedegreen Research — Where Did the Saved Hours Go? (https://hedegreenresearch.com/articles/the-saved-hours-doctrine/). Conceptual question about the allocation of saved labour time, not a measured saving in this essay.

[13] Hedegreen Research — Why Should Anyone Go Hungry? (https://hedegreenresearch.com/articles/why-should-anyone-go-hungry/). Food-floor research questions.

[14] Hedegreen Research — Money Should Not Be the Door to Survival (https://hedegreenresearch.com/articles/money-should-not-be-the-door-to-survival/). Human Baseline conceptual lineage.

RELATION MEMORY

CONTINUES I Do Not Want To Lose This · Part II revisits the original access and production-capacity question.

RETURNS TO The Thinking Machine · The March reflection predates the terminology used here; this essay gives it an operational use.

NEARBY What If the Slop Is Mine? · Production and verification debt are separate questions.

NEARBY Money Should Not Be the Door to Survival · Conceptual lineage for Human Baseline, not independent evidence of feasibility.

SOURCE LAYER The World Kept Checking · Concrete research example; captured evidence records the work and its limits.

NEARBY Where Did the Saved Hours Go? · Who receives the benefit of saved labour time is a separate question from increased output.

CANONICAL URL: HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/I-STILL-DONT-WANT-TO-LOSE-THIS/