

**Pro Tools for
Superintelligence.**
Extraordinary intelligence.
In individual hands.

SuperDuck.AI

Duck powered.
Your thoughts.
SuperDuck's cognitive muscle.

THE PERSONAL AI EDITION

OCTOBER 1, 2026 · IMAGINED EDITION

NO. 01 · FREE TO READ

THE PERSONAL COMPUTER HAS NOT FINISHED ITS WORK.

SUPERINTELLIGENCE / THE SUPERDUCK VIEW

ARTIFICIAL DINOSAURS BUILT! SCIENTISTS WARN OF EXTINCTION!

Experts demand emergency powers to save mankind!

How a failure of imagination turns a technology of human empowerment into a story about a rival species.

THE SUPERDUCK VIEW · CAPABILITY

Suppose hydraulics had been invented in our own intellectual climate. An engineer discovers that fluid under pressure will push harder than any ox. One man at a lever can now lift a locomotive, hollow out a hillside, and shift in a morning what a hundred laborers shifted in a month.

You would expect talk of bridges and tunnels. Instead a committee of influential persons decides that the true significance of the invention is a new race of powerful beings. They dress the excavators in plastic hide, bolt on teeth and a tail, and announce the Artificial Dinosaur. AD has arrived.

The creatures are stronger than any man. They can pull down a building. They can flatten a carriage. Some can be left running with nobody at the lever. Plainly we have made a new species, and the questions follow of their own accord. What happens when the Dinosaurs notice they are stronger than we are? Will they obey us forever? Might they form plans of their own? Ought ordinary people to be allowed to drive one?

Institutes are endowed. Conferences convene. A profession springs up devoted to the existential risk of the Artificial Dinosaur, and its leading members conclude, after study, that only a few approved institutions can be trusted with the beasts. Meanwhile a man who understands hydraulics is down at the river trying to build a bridge.

The mistake is in the metaphor

When machines grew stronger than men, nobody declared them a superior order of being. We said that men had got hold of a great deal of power. The excavator was never the hero of the industrial age; the fellow in the cab was, and a man who once dug a cellar with a spade could now dig a harbor.

Superintelligence wants the same treatment. We are building ma-

chines that reason, learn, analyze and design, in places better than we do unaided. That is a wonder. But the wonder need not be a rival species. It can be a human being with a very large mind at his disposal, which is a different thing and a better one.

From artificial muscles to artificial intelligence

Hydraulics gave a man forces no muscle could supply. The computer gave him arithmetic no brain could keep up with. Superintelligence gives him depth and speed of thought that unaided attention cannot sustain. A hydraulic cylinder has never been caught telling a lie, but it can be poorly controlled. An intelligent machine may give a wrong answer with perfect confidence. That calls for care, certainly. It does not call for turning the machine over to a committee. Do we organize our thinking around what the machine is supposed to want, or around what its owner can now do?

Keep the human in the story

Hydraulics did not give us artificial animals. It gave us better ways to move the earth. Superintelligence is no different and should not be used to erect a new authority. It should be used so that people can understand, discover, invent and build more than they could before. Take the risks seriously; build safeguards and harnesses to match the capability. Keep a clear line between a tool that assists a decision and a system that acts with nobody at the lever. But do not let a mythology about machines crowd out the people the machines are for.

The future belongs to the individuals who can at last build what they have only been able to imagine. It should not belong to fear-mongers of Artificial Dinosaurs, nor to the people who wish to appoint themselves their keepers.



SuperDuck works with you.

Research, reasoning and creation on the user's own machine.

*Your curiosity, ideas, and taste.
SuperDuck does the work.*

INSIDE THIS EDITION

The individual: learning without permission	PAGE 2
The priests, the police referral, the confidential data	PAGE 3
Freedom to innovate, how it works & the Academy	BACK PAGE

SUPERINTELLIGENCE / LEARNING

From answers to understanding

A calculator will give you the answer without teaching you arithmetic. A search engine will find an explanation without troubling itself over whether you understood a word of it. Both are useful machines, but neither is much of a teacher.

A teacher worth the name does more than supply answers. He finds out where the pupil went astray. He asks a question, offers an example, and, if necessary, begins again by another road.

An intelligent machine can do much the same. Suppose a man wishes to understand electricity. He may begin with a question about a battery and discover that he has never properly understood a circuit. The machine can take him back to first principles, follow his questions, and proceed as quickly or slowly as his understanding permits.

The remarkable thing is not that the machine knows the answer. It is that the man need no longer pretend to understand it.

CONTINUED ON PAGE 2

PERSONAL COMPUTING / THE SUPERDUCK VIEW

SuperDuck: a workshop for the mind

Extraordinary intelligence. In individual hands.

A carpenter has his tools, a machinist his lathe, and a musician his instruments. A person engaged in intellectual work ought to have a workshop of his own, too.

SuperDuck is such a workshop. It puts artificial intelligence to work on the owner's computer, with the owner deciding which intelligence to employ. One may use a model running on the Mac, Apple's own intelligence, or call upon more powerful machinery elsewhere. Toasted AI, supplied through LinuxToaster, is available when wanted. Other models may be reached with one's own API keys. There is no requirement to subscribe to a particular intelligence or remain faithful to a single one.

The workbench is a folder of ordinary files. Bring in a document, examine an argument, revise a manuscript, or ask the machine to find an error in a program. Commands can refine, shorten, trans-

HOW IT WORKS, BACK PAGE

late, check and improve the work. An instruction can be repeated until it meets a stated test. And when the machine claims to have changed a file, SuperDuck can show a receipt for the deed, not merely its promise.

There are instruments for pictures and sound as well. The machine can examine an image, draw diagrams in its own graphics language, generate pictures locally, take dictation and speak. A job may continue while its owner attends to another.

The remarkable thing is not that every tool is new. It is that so many tools can now be placed at the disposal of one person, working at his own pace, with his own materials, on machinery he controls.

We call it **Pro Tools for Superintelligence**. The purpose is not to have the machine supply the ideas. It is to give the person who has them a better chance of carrying them out.

A LITTLE HELP IS AT HAND

Get SuperDuck.

Start with one document, one question or one idea. Free beta; Apple silicon Mac required.



Scan the code or visit superduck.ai

ANALYSIS

READ THE EVIDENCE. DECIDE FOR YOURSELF.

COMPANIES / ORGANIZATION

Companies organize people; the minds stay their own

Their purpose is not to own the intelligence of their employees. It is to organize unreliable individuals into reliable systems.

THE SUPERDUCK VIEW · CAPABILITY

A company's purpose is not to own the intelligence of its employees. Its great achievement is to organize unreliable individuals into a reliable system. People are inconsistent; they differ in ability and attention, they fall ill, they make mistakes, they quarrel and they change their minds, and a well-run company takes this material and produces, every day, the same product to the same standard. It does so by coordination, process, accountability, cap-

ital, shared purpose and the habit of continuing.

Superintelligence does not abolish that function. It raises its value. When every individual becomes dramatically more capable, the hard problem is organizing all that capability into something that can be relied upon. The best companies will not be the ones that use superintelligence to get by with the fewest people. They will be the ones that find out how much can be done with the people they have.

SUPERINTELLIGENCE / THE SUPERDUCK VIEW

Boardrooms mistake AI for outsourcing

The great mistake of corporate AI strategy is confusing the elimination of work with the expansion of human capability.

THE SUPERDUCK VIEW · CAPABILITY

The personal computer did not transform business because executives discovered they could fire the accountants. It transformed business because the accountants could suddenly ask questions they had never been able to afford. Before the spreadsheet, every additional question cost a day of arithmetic; a second scenario meant a second day, and curiosity stopped where the arithmetic did. Afterward an

accountant could try a dozen scenarios in the time one had taken, and questions that had never justified the labor became routine. He did not become obsolete. He became more capable, and that is how the personal computer got into the company: as an instrument of personal augmentation first, corporate infrastructure second.

CONTINUED ON PAGE 2

SUPERINTELLIGENCE / THE SUPERDUCK VIEW

AI exposes work that no longer needs doing

There is a deeper problem with the labor-replacement story. It assumes the work people do is necessary. A good deal of it is not. Much of what goes on in an organization exists to make up for limits on what people can process, remember, calculate and coordinate. The report is written because the information is hard to get at. The meeting is held because coordination is hard. The administrative process exists because two systems cannot understand each other. All of it is necessary under present constraints, and none of it is valuable in itself.

The object was never the report; it was to know what was happening. It was never to reconcile the spreadsheets; it was to have an accurate picture. It was never the meeting; it was the decision. When superintelligence removes the constraint, the intermediate activity can simply stop, which is not the same as replacing the person who performed it. The person is now free to go after the object directly.

Work is not the objective. Human achievement is the objective. Work is what we have had to do to get there.

Your thoughts. SuperDuck's cognitive muscle.

The greatest product is the individual

For the first time, the speed of learning can be organized around the learner rather than the institution.

THE SUPERDUCK VIEW · LEARNING

Human potential has always been limited by access: to knowledge, to teachers, to mentors, to institutions, to anyone willing and able to explain the thing you do not yet understand. A brilliant student with a poor teacher is at a disadvantage. An entrepreneur without an experienced adviser learns by expensive mistakes. An adult entering an unfamiliar field works through materials written for somebody else’s background and pace. We built enormous institutions to ease these problems, and every one of them has the same limitation: it must organize learning for the group.

Superintelligence changes the economics of learning. A person can have a patient, knowledgeable, adaptable intellectual partner on hand whenever she wants to learn something.

Learning at the speed of the individual

Consider a student who meets a hard idea in mathematics. In a classroom the teacher sets a pace for the room; some understand at once, some need more time, some need a different ex-

planation altogether, and the lesson moves on regardless. With a capable AI mentor she can ask the same question twenty times without anyone sighing, ask for another explanation, an analogy, a worked example, or to run the problem backward from the answer. The mentor can find out what she does not understand and build the missing foundation. The quick student need not wait and the slow one need not fall behind.

The individual no longer learns at the institution’s speed; the institution will have to learn to keep up with the individual. The scarce thing in education was never information. The library has been there for centuries and the internet put most of it in a pocket. The scarce thing was personal intellectual attention, and that is what superintelligence can make plentiful.

The compounding effect of capability

Learning compounds. A new mathematical technique opens physics; physics opens engineering; engineering lets a person design what she could only picture before. Each capability is

a key to the next, and the rate at which the keys are collected has been limited by time, money, mentors and admission.

Consider someone who wants to build a medical device, start a manufacturing company or understand an unfamiliar science. Until lately that ambition meant years of formal study before she could ask a useful question. With an intelligent mentor she can begin at once, learn the foundations as she needs them, test what she thinks she knows and find out where deeper expertise is required. The machine does not spare her the experiment, the practice or the judgment, and it does not guarantee that its own advice is right. What it does is put the intellectual tools within reach, and the point of learning faster is not to arrive sooner at the old destination. It is to set out for the ones that were out of the question.

The democratization of mentorship

None of this should depend on working for a large corporation, attending an elite university or living near

a center of expertise. A teenager in a small town should be able to explore advanced mathematics. An electrician should be able to learn electronics engineering, a nurse statistics and research methods, a retired welder programming or anything else he fancies for the next thirty years. They do not need an institution to decide in advance which ambitions are appropriate to them. They need tools that help them pursue the ambitions intelligently.

That is why superintelligence is first of all a personal technology. Its value is not that a company can buy cognitive labor more cheaply. Its value is that an individual can become more capable than her circumstances had allowed.

The most important product of superintelligence is not the work it performs. It is the individual who becomes more capable through using it.

What people can now do

We are used to measuring a machine by the labor it saves. So many hours spared, so many tasks completed. That is a modest way to measure an invention that may change what people can attempt in the first place.

Consider a man who hears music in his head but cannot play the piano. A woman who can picture a scene but has never mastered a brush. An inventor who knows what a machine ought to do but cannot draw its mechanism. Until now, each faced a long apprenticeship before the first experiment could begin.

Years of exercises, technical manuals and tiresome mistakes stood between the idea and its expression. The new machines can perform much of that preliminary labor. A person may begin with an idea, examine the result, correct it and try again. Imagination need no longer wait upon every mechanical accomplishment.

A child has questions enough to exhaust a schoolhouse. Now there may be a patient explainer at hand, willing to begin again, put a difficult subject another way,

and answer the twentieth question as readily as the first. The young need not know everything before they can begin to understand something.

An older person may have a lifetime of knowledge but find that memory, eyesight, dexterity or stamina no longer serve as faithfully as they once did. An intelligent assistant can help retrieve a forgotten detail, read difficult material, organize a complicated undertaking or turn spoken ideas into written work. What time takes from the machinery of a person, another machine may help supply.

None of this abolishes the need for skill or judgment. A machine may produce the notes, but somebody must decide whether the music is worth hearing. It may assist an old craftsman whose hands have grown unsteady, but it cannot manufacture the experience those hands acquired.

The question is no longer how much work a machine can save. It is how many more people can now make something worth doing.

CONTINUED FROM PAGE 1

Boardrooms mistake AI for outsourcing

Much of corporate leadership approaches superintelligence as the next generation of outsourcing. The reasoning is tidy. Companies pay people to do cognitive work; AI does cognitive work; therefore AI lets companies pay fewer people. It is the least imaginative use that could be found for one of the most consequential inven-

tions ever made. Outsourcing moved existing work to a cheaper address. Superintelligence lets an individual work at a level that was closed to him: investigate an unfamiliar field, write software, model a problem, test a hypothesis, without first assembling a department of specialists. The question is not how cheaply the old work

can be reproduced. It is what becomes possible when a person is no longer confined to the size of his own head.

A company that looks at superintelligence and sees headcount is a contractor who, shown the first hydraulic excavator, asks how many shovel men he can let go. He has noticed the least in-

teresting thing about it.

The wrong question

Cost reduction has a floor; eventually there is nothing left to cut. Human capability has no corresponding ceiling. People who learn faster, explore more and solve harder problems create value that did not exist before, and

that is the opportunity the headcount view misses. The spreadsheet did not merely make accounting cheaper; it enlarged what financial reasoning was. The personal computer did not merely make clerical work cheaper; it made computing a personal activity and created industries. Superintelligence is the next such enlargement, and proba-

bly the largest. Companies should stop asking how many people they can replace and start asking what their people can now become capable of.

Superintelligence is for people. The company is the system that organizes what those people can accomplish together.

CONTINUED FROM PAGE 3

Cloud AI turns confidential data into evidence

The woman in Bonita Springs (page 3) used Claude as a diary. She wrote a threatening entry, Anthropic’s safety process flagged it, and law enforcement became involved. She said she was venting; investigators saw a threat. Whatever a court makes of it, she did not contact the police. The system did. That is placement: the information crossed a boundary she no longer controlled.

Layering

Feed a cloud model your notes, your mail and your research and ask for a summary. What comes back is cleaner.

Your voice is gone, your typos are gone, your three-in-the-morning tone is gone, and the result reads like a neutral account of what began as a mess. Ask where it came from and the answer is “an AI”, which tells you nothing about the material behind it. The trail has not necessarily vanished; logs and records may exist. But the output does not carry its history.

The same applies to a proprietary codebase, an unpublished manuscript or a trade-secret process sent out for improvement: the model may not have learned from it, but somebody else has

received information you were responsible for protecting. It applies to reputation too. A selective pile of someone’s private messages comes back as a seemingly neutral summary of his views; the words survive and the circumstances disappear.

Integration

Integration is when the polished result re-enters the world as ordinary work: a report, a pitch, a legal filing, a story. The reader sees the finished product, not the private documents behind it, and nobody needed to intend anything. The person wanted a summary, the

machine supplied one, it was useful, so it went in the report. The output looks cleaner and its origins are harder to trace. The obligations attached to the original have not gone anywhere.

The anti-laundering feature

SuperDuck is built around this distinction. Run the model on your own Mac, work on files in your own folders, keep the results there. If the information never leaves your computer there is no placement; if it is processed locally there is no outside party doing the transformation; if the output stays in

your files there is no integration into a system you do not control. The optional browsing and external computation move that boundary, which is why they are switches.

A local machine is not invulnerable: backups, synced folders, malware and remote access can all carry information out, and running locally does not by itself discharge a professional obligation. What it does is remove a whole category of third-party exposure. The cloud turns private inputs into clean-looking outputs whose origins and obligations are easy to overlook. Call it

information laundering.

If your job involves non-public information, your AI should be on your machine. Not because you have something to hide. Because you have something to protect.

An argument, not legal advice. The referral case is from the public record and is set out on page 3; the pharmaceutical example is hypothetical.

CONTINUED FROM PAGE 1

From answers to understanding

He may ask the foolish question, the elementary question, or the same question for the twentieth time. The machine will not sigh, look at its watch,

or remind him that the rest of the class has moved on.

Nor must the machine always hand over the answer. Sometimes the bet-

ter service is to ask the next question, point out a contradiction, or let the pupil work through the difficulty himself.

There is a danger, of course, in accepting every explanation merely because it is delivered with confidence. The machine can be mistaken. The pupil must

still learn to test what he is told.

The object is not to make a man dependent upon a machine that knows everything. It is to leave

him knowing more than he did before, and better able to find out the rest for himself.

CONTINUED FROM PAGE 4

SuperDuck opens its workshop

Bring an idea. Leave with something made.

A mechanic may arrive with a machine he wishes to improve. An accountant may bring a question that has never justified a week of calculation. A writer may bring an unfinished manuscript. A curious person may simply bring a subject he has wanted to understand for twenty years.

The seminar begins with what the par-

ticipant wishes to accomplish, not with a list of things the software wishes to demonstrate.

There will be mistakes, revisions, discarded attempts and, one hopes, a few surprises. That is how workshops have always operated.

The object is to leave with something made, something understood, and enough command of the instru-

ments to continue without the instructor.

The tools come with the instruction

SuperDuck software and course materials are included in the seminar fee. Optional outside computing may involve additional usage charges, which will be explained before the machinery is engaged. Participants may also use their own API credentials.

The workshops will be open both to private individuals and to companies wishing to enlarge the abilities of their employees. Business participants will work with approved materials; the lessons belong to the people who learn them.

A company has good reason to want capable employees. But the employee’s mind is not a piece of company equipment, and an education worth having

ought to remain useful after the working day is done.

There is a difference between learning to operate a machine and learning what one can accomplish with it. SuperDuck intends to teach both.

A carpenter does not become a carpenter by attending a lecture on the history of hammers. He becomes one by building something.

The same principle ought to serve for superintelligence.

SUPERDUCK ACADEMY · Practical instruction in personal superintelligence. Public workshops and private company sessions are being prepared. Details and enrollment information: **superduck.ai**.

Where are the intellectual skyscrapers?

What are the bridges of superintelligence? A farm boy with a better tutor in physics than any university can hire. An inventor who can try an engineering idea without first raising the money for an engineering department. Three people in a rented room doing the work that once took a research institute. A pharmacist who can read the trial her customers ask about. A clerk learning a trade at the speed of his

own curiosity rather than the speed of the semester. A retired machinist who finally understands the quantum mechanics he gave up on in 1971.

Those are the intellectual skyscrapers. A civilization that sees only a dangerous new animal will not build them, and the loss will not appear in any safety report, because nobody counts the bridges that were never built.

The personal computer enabled individuals; AI follows

For most of its early history computing lived in large institutions. The mainframe was expensive, time on it was scarce, and an operator with a clipboard decided whose job ran next; that was partly the machinery’s fault and partly the economics. The personal computer ended the arrangement. A person could program, calculate, experiment and build without applying to the computing department, and the computer became personal.

The price of arithmetic fell, but the larger event was a redistribution of agency. People found uses for computers that the institutions had never anticipated; software companies started in bedrooms and garages; experiment went decentralized. The large organizations did not disappear. They lost the one justification they had enjoyed for free, control of the machine, and had to find value elsewhere. Superintelligence is about to do the same to cognitive capability.

Safety rules should arm the individual

Be careful in proportion. Hold the maker liable for harm. Test a system before it runs a power grid, and do not test the one that tutors physics as if it did. Pay for independent evaluation. Draw the line where it belongs, between a tool that helps a person decide and a machine authorized to act with

nobody’s say-so, and regulate the second. Give people the knowledge and the transparency to control what they use.

A good safety regime leaves people better able to use a technology. An extractive one leaves them more dependent on whoever holds the keys.

Fear of AI pays

An excavator can knock down a house, cut a gas main, or roll over on its driver, and the world’s answer has been licenses, training, guards on the moving parts and a man who is liable when it goes wrong. Nobody proposed that excavators be owned only by a few insti-

tutions with a special understanding of mechanical force.

Some proposals for governing AI propose exactly that, and the burden belongs on the people asking for it to show why a license and a liable owner would not do.

SUPERINTELLIGENCE / ACCESS

Priests of AI guard the gate

Every great democratization of capability challenges the institutions that once controlled access to it.

THE SUPERDUCK VIEW · ACCESS

Whenever a powerful capability comes within reach of ordinary people, the same argument is made against them. It is too important to be left to individuals. It is too complicated for them to understand. It is too dangerous to use without supervision. Therefore a properly qualified institution must stand between the people and the thing.

The institution varies. The church once held the exclusive right to read the scripture to the congregation. The computing department held the keys to the mainframe and the schedule of who might run a job on it. The guilds decided who could call himself a surgeon or a lawyer. The logic does not vary at all.

Now superintelligence proposes to put extraordinary cognitive power in individual hands, and the argument is back in its old clothes. The technology is too powerful. The public cannot be trusted with it. Only approved organizations should work at the frontier. A priesthood of intelligence is forming, and it would like us to come to it for our thoughts.

The economics of the intermediary

An intermediary can be worth having. A teacher, a physician, an engineer: expertise matters. But there is a difference between offering expertise and requiring everyone to obtain permission from the people who hold it. An institution that controls access to a capability acquires power over everyone who needs it, and with the power come fees, influence and prestige, and an interest in keeping the arrangement that produces them.

The simplest way to keep it is to dwell on the dangers of unmediated access.

The method works best when the dangers are partly real, and they usually are: complicated tools are misused, novices make mistakes, specialists get better results. None of that establishes that one institution should hold the keys. The question to ask of any intermediary is whether it is making people capable or keeping them dependent.

Intelligence without permission

Picture a person who can get expert help with mathematics, programming, engineering, economics, writing or research by asking for it. She does not enroll anywhere first. She does not need her employer's leave to be curious, or a licensed gatekeeper's opinion of whether her ambition suits her station. She begins, experiments, works out her own understanding and follows ideas that would otherwise have stayed in the drawer.

That is not a world without experts or institutions. It is a world in which they must earn their place by helping people accomplish things, rather than by controlling the supply of intelligence, and some of them will not like it.

Fear, uncertainty, doubt, and regulatory capture

Fear is a serviceable instrument of institutional preservation. When a technology is new and powerful the public cannot easily judge claims about its risks, and an organization that presents itself as the indispensable protector will find an audience. Some risks deserve serious governance: systems that act on their own, that can be turned to breaking into other systems, that deceive, that fail at scale. But the existence of a risk is not a

blank check.

A regulatory regime can protect the public and protect incumbents in the same motion. Compliance that only the largest companies can afford thins out the competition. Licensing that has little to do with safety still keeps the newcomers out. Restrictions on access keep people from acquiring the experience that would let them use the thing responsibly. The danger is that safety becomes the argument for scarcity, and scarcity is the one thing superintelligence can abolish.

Who owns the future of intelligence?

The political question of superintelligence is access. Will extraordinary cognitive capability be something individuals use to develop themselves and pursue their ambitions, or a service dispensed by a few institutions that decide what people may know, make and attempt? The first path is the personal computer's. The second rebuilds the mainframe with intelligence itself as the rationed resource.

There will always be work for institutions; we need reliable systems, common standards and someone to hold accountable. But an institution should serve human capability, not claim title to it.

Every great democratization of capability challenges the institutions that once controlled access to it. Superintelligence may be the largest such democratization yet. The question is whether we let it become personal, or let a new priesthood persuade us that intelligence, like salvation, must come through them.

SPECIAL REPORT / PRIVACY, POLICY AND THE PERSONAL COMPUTER

When a private chat reaches the police

In defense of unfinished thoughts. As AI companies begin referring users to the police, an old distinction acquires new weight: a computer that works for its owner, or a service that also watches what its customers say.

BY OUR TECHNOLOGY CORRESPONDENT

Someone opens an AI assistant late in the evening and types something they would not say aloud: an angry thought about a colleague, an argument they are testing, a scene for a novel, a question about how a crime is committed. The words are clumsy and half-formed, and the person is often trying to find out what they think rather than announcing what they mean to do. The question this report asks is who else can read them.

For most of the history of the personal computer the answer was nobody. A document written on your own machine stayed there unless you sent it somewhere or someone took the machine. With a cloud chatbot the conversation happens on another company's computers, under rules that company writes and can rewrite, and those rules now say when it may call the police.

When a conversation becomes a referral

In August 2025 OpenAI set out its procedure: conversations that suggest a plan to harm others are routed to a small review team, which may ban the account and, where it finds "an imminent threat of serious physical harm to others", refer the case to law enforcement. Anthropic's privacy policy permits disclosure where the company believes it "reasonably necessary to ... prevent serious harm to any person or to property".

These are not dead letters. After the school shooting in Tumbler Ridge, British Columbia, in February 2026, in which eight people were killed, OpenAI disclosed that it had banned an account belonging to the attacker the previous June, had considered reporting him and had decided the threshold was not met. Sam Altman, its chief executive, wrote to the town that he was "deeply sorry". The company told Canadian ministers it had since made the threshold "more flexible" and would refer such an account "if it were discovered today".

Other cases went the other way. Darren Shida Zhou, a former Goldman Sachs analyst in Florida, told ChatGPT repeatedly, and in detail, that he would kill his former girlfriend. OpenAI reported

him to the FBI, and in August 2026 he was sentenced to eight years' probation. At least three Claude conversations reached the police in August and September 2026. One belonged to a woman in Bonita Springs who, by her own account, used the chatbot as a diary. She wrote that she would "shoot up" the county sheriff's office; four days later she was charged with a felony. What people tell an AI, the sheriff said, is "never truly anonymous".

Nobody argues that a company holding evidence of an imminent attack should look away, and the Zhou case shows why. The harder questions are the ones the policies leave open: who gets to read a person's unfinished thoughts, who decides when a thought has become suspicious, and whether a private company should be the one making that call.

Thinking out loud

Thinking is untidy. People try out explanations before they accept them, overstate a case to see if it holds, describe things they would never do and say things they take back. An assistant that is any use has to work with all of that.

A business owner who types, after a disastrous project, that a client deserves to be ruined is angry, and a good assistant asks what happened and helps her work out what to do about it. The sentence says nothing about what she intends. The same goes for a novelist working on a villain, a reporter reading extremist forums or someone trying to make sense of a thought that frightens them.

Physics and policy

Privacy should not depend on a provider's restraint. When a model runs on a personal computer with networking and external processing off, the conversation never travels to an inference provider, and a company cannot disclose what it never received. That is a more durable promise than any privacy policy, and it is the one SuperDuck.AI makes for offline processing.

There is a harder way to put it. Money laundering places cash in the system, layers it until its origin is lost and brings it out looking clean; a cloud model can reproduce a similar pattern with what-

ever is typed into it, and the person who typed it carries the risk. The argument is set out on page 3 and continued on page 2.

A computer at home is still subject to the law. The police can seize it with a warrant, and synced folders, backups, malware and any online service you switch on are other ways for its contents to leave. Nor does running locally make a model fair, accurate or free of restrictions. The point is narrower: where the processing happens decides who can watch it.

A room of one's own

SuperDuck's arrangement is simple. You bring the purpose, the material and the judgment; the software brings the cognitive muscle. With a model installed and external services off, conversations and document work happen on your Mac and no cloud provider is involved. The optional connections change that, and are worth understanding before you switch them on.

A good assistant will still argue with you, pick holes in a weak case and point out what you have not considered. What it need not do, running locally, is send the conversation to a company that may also be judging you.

AI companies will go on facing hard cases, and no rule settles all of them. What people should have is a choice, rather than a provider's policy as the only protection on offer. Some will stay with cloud services, some will keep a model on their own hardware, and many will do both. A personal computer should be a place to start with a doubtful idea, test it, find it wrong and try again, without anybody else in the room.

The room to think belongs to the person doing the thinking.

The cases above are from the public record; the Lee County charge is an allegation, and the defendant had not been arraigned at the time of writing. Local inference prevents disclosure to an inference provider only where the content is not transmitted through other services; users should verify their configuration, device security and legal obligations.

THE RECORD

How the rules changed in eighteen months

MAY 13, 2025

A federal judge orders OpenAI to preserve every output log on its consumer plans, including chats users delete, for the New York Times's copyright suit. Enterprise and zero-retention API customers are excluded.

JUNE 2025

OpenAI bans a ChatGPT account for violent content and decides a police referral is not warranted. The account holder carries out the Tumbler Ridge shooting eight months later.

JULY 2025

Sam Altman says ChatGPT conversations have no legal privilege: "we could be required to produce that."

AUG. 27, 2025

OpenAI publishes its procedure: harm-to-others conversations go to human review and "may" be referred to law enforcement. Self-harm cases are not referred.

OCT. 9, 2025

Dates from company statements and court records as reported; see sources.

The court ends the preservation duty for new chats, effective Sept. 26. Logs saved before that date stay under hold.

FEB. 26, 2026

After Tumbler Ridge, OpenAI tells Canadian ministers its referral threshold is now "more flexible" and the banned account would be referred "if it were discovered today".

AUG. 2026

Darren Zhou, reported to the FBI by OpenAI, is sentenced to eight years' probation. Two Claude conversations reach police in the same month.

SEPT. 10, 2026

Anthropic's privacy policy takes effect with disclosure permitted where "reasonably necessary to ... prevent serious harm to any person or to property".

SEPT. 30, 2026

A Florida woman who used Claude as a diary is charged under the state's written-threats statute after Anthropic's reviewers report an entry.

SPECIAL REPORT / PRIVACY

A machine that can watch holds no warrant to report

We have been here before. A word processor does not read every sentence for signs of criminal intent and a spreadsheet does not check its cells for fraud, though both could be built to and both have been used to plan and hide terrible things. The companies that make them do not report their customers. Nobody finds this strange, because the expectation was settled decades ago: a person who buys a machine and pays for the software on it may use them in private. The tool helps its owner write a sentence or work a sum. It does not pass judgment on either.

The law takes the same view. Being able to get at information has never been reason enough to take it. The Fourth Amendment limits government searches. Federal law restricts what providers may disclose from stored communications, with an exception for danger of death or serious injury, which is the provision under which a company may volunteer a chat to the police. The Supreme Court has said that digital records expose the intimate details of a life. The protections are patchy

and the rules for companies are looser than those for the police, but they rest on the same idea as the word processor: you do not give up your privacy because somebody else has the technical means to look.

AI muddies this because an assistant does more than record words. It can advise, reason and help someone act, and there are good reasons for it to refuse to help with serious wrongdoing. Refusing is one thing; reporting the person who asked is another, with a different weight of responsibility behind it. A model that can spot a troubling thought has not thereby earned the right to investigate its author. If preventing harm were a general license to inspect private thinking, the license would cover every document, spreadsheet and half-finished idea a computer could read, and personal computing has never worked that way. SuperDuck belongs to the older tradition: software bought to run on the owner's machine, working for the owner and in no position to tell anyone what it saw.

SPECIAL REPORT / PRIVACY

Bad first drafts deserve privacy

Everyone who has kept a notebook knows what it is for. You write down the argument you are not sure of, the letter you will not send, the plan that looks foolish by morning, and you find out what you think by seeing it on the page and crossing most of it out. The crossing out is the useful part, and it is the part an assistant can help with, by asking what the argument rests on or raising the objection you had not thought of.

COMPANIES / NON-PUBLIC INFORMATION

“Nothing to hide” is no defense for a leaked secret

“I have nothing to hide” sounds reasonable until you count how much information does not belong exclusively to the person holding it. A lawyer has client communications, a physician patient records, an engineer proprietary designs, a researcher unpublished re-

Nobody writes well with a stranger reading over the shoulder. The sentences get careful, then they get safe, then they stop. A question that might be reviewed by someone at a company does not get asked, or gets asked in the form most likely to pass. That is not an argument for ignoring a credible threat. It is the plain observation that watching changes what is watched, even when the watching is well meant, and that the first draft, the bad one, is where the thinking happens.

The question is not whether you have something embarrassing to hide. It is whether you have something you are responsible for protecting.

COMPANIES / NON-PUBLIC INFORMATION

Cloud AI turns confidential data into evidence

When you handle information that isn't yours to disclose, convenience is not a defense. Cloud AI introduces risks that good intentions cannot erase.

THE SUPERDUCK VIEW · COMPLIANCE

Confidentiality is not a promise to keep quiet. It is an obligation to control where information goes. Money laundering has three stages, placement, layering and integration, and each has a parallel in what happens when private information is fed to a cloud AI.

Placement

Placement is the moment raw material enters the system: the diary entry, the angry draft, the half-formed business idea. Typed into a cloud service, it sits on someone else's computers, where it may be logged, reviewed or processed under arrangements you do not fully understand.

COMPANIES / CASE STUDY

An unannounced drug trial meets the cloud

Consider a pharmaceutical researcher who types into a cloud model: “The Phase 3 results are better than expected. The announcement is Thursday.” Perhaps she only wanted help with a presentation. Before Thursday that sentence may be worth a great deal; after Thursday, when the results are public, most of its trading value is gone.

An invention is worth protecting because of what it is. Market-sensitive information is worth protecting because of when it becomes known. The model does not need to decide to trade. The informa-

tion only needs to reach someone who might, and depending on the service it may be retained, or accessible to staff and systems, or exposed by a breach before the announcement.

Whether sending it to an AI service breaches a legal duty depends on the law, the contract, the safeguards and the permissions involved; a governed enterprise service is not a consumer chatbot, and not every disclosure is a crime. But the researcher may still have to answer a question with no good answer: why was it sent there at all?

PRIVACY / HOW IT WORKS

Veil sends the question and keeps the names

Sometimes the most capable model is not the one on your Mac. Veil is for those occasions. Before material leaves the machine, it substitutes placeholders for the particulars that would identify the subject: the names of people and organizations, the places, the dates and figures. The outside model works on the veiled text and returns its answer, and SuperDuck restores the real particulars on your own computer. The outside intelligence did the thinking and learned less about whom it was thinking for.

That is a better bargain than a promise of privacy. A promise depends on the other party's conduct. A veil depends on

what you sent.

It is not a disguise that cannot be seen through. A memo about a drug company in “[City]” whose trial of “[Drug]” failed in “[Month]” may still be identifiable to anyone who knows the field; the surrounding context can name what the placeholders hide. Veil reduces what an outside model learns. It does not guarantee anonymity, and the person who knows the subject best is the one to judge how much the rest of the text gives away. For material that must not leave the machine at all, the local model remains the answer.

CONTINUED ON PAGE 2

OPINION / THE SUPERDUCK VIEW

Freedom to innovate

The promise of artificial intelligence is not merely safer machines. It is a world in which more people have the freedom to learn, think, experiment, invent and create.

The history of invention is crowded with people who were told that their ideas would not work. A good many of those ideas did not work. A few changed the world, and the difficulty has always been that nobody could tell which was which until somebody tried.

Trying is expensive.

A curious person may need a mathematician, a machinist, a lawyer, a translator, a patient teacher and three months in a library before an idea can be tested. Most people cannot assemble such a company, so most ideas are never tried at all.

Now imagine that the curious person can begin with the computer on the desk.

A schoolteacher might work through the engineering of a device she has been thinking about for years. A mechanic might learn the mathematics needed to improve a machine. A small manufacturer might investigate a process that previously required a research department. None is guaranteed success. That is not the point.

The point is that they can begin.

We have grown accustomed to measuring technology by the labor it saves. We should also measure it by the ambitions it makes possible. The first calculation concerns the work we already know how to do. The second concerns work that nobody has yet attempted.

SuperDuck is built for the second calculation.

The purpose of personal superintelligence is not to make every person an expert in everything. It is to give more people a fair chance to become expert in something, and to put their knowledge to work without first obtaining permission from a large organization.

A machine cannot supply courage, judgment or an idea worth pursuing. But it may remove a surprising number of the obstacles that stand between an idea and its first experiment.

That is a considerable service to civilization.

The publisher's argument. SuperDuck.AI develops personal AI software for the Mac.

HOW IT WORKS / PHYSICS VS. POLICY

Your Mac supplies the muscle

The Mac supplies the computing power. Here is how local inference works, what it enables and what still needs checking.

TECHNICAL GUIDE

SuperDuck brings powerful AI capabilities to the Mac you already own. The intelligence runs locally, using the computing power of Apple silicon, rather than requiring every question to travel to a remote server.

SuperDuck turns that intelligence into a personal assistant for researching, reasoning, writing, analyzing information and developing ideas. You can also access it from your iPhone or iPad.

The important distinction is who controls the computer, the information and the purpose of the work.

With SuperDuck, that person is you.

Where the work happens

With a model installed, chat and document work run on the Mac. The model, the conversation and the files stay on that machine; after setup the Mac can be disconnected from the network and local chat

keeps working.

Remote access changes one thing. The iPhone or iPad sends the question to the Mac over an encrypted link and the Mac sends the answer back. The processing still happens on the Mac, which has to be on and connected, and the conversation now travels between two devices you own.

The optional services change more. Browsing sends requests to websites, accelerated inference sends the conversation to external hardware, and model downloads reach download servers. Each is off until you switch it on; the table below sets out what each one sends and where.

Your computer, your control

Files in a synced folder may have remote copies even when inference stays local. Device access, encryption and backup settings decide what happens around the process-

ing. The eight questions at superduck.ai cover the full setup.

Anyone handling regulated or client-confidential material should check the whole setup against their obligations before relying on it. Local inference settles where the processing happens and nothing else.

No account, no telemetry, no update checker. SuperDuck does not require an account to run, sends no usage data and does not check for updates on its own. Nothing leaves the machine that its owner did not send.

What the firm knows stays on the firm's machine.

TEST ONE CLAIM AT A TIME

Check the claim yourself

A simple offline check shows whether local chat depends on a cloud service.

- 1

FINISH SETUP.
Install SuperDuck and complete the model download. Start with a non-sensitive sample document.
- 2

SELECT LOCAL PROCESSING.
- 3

DISCONNECT THE MAC.
Turn off Wi-Fi and unplug Ethernet
- 4

DO A FRESH PIECE OF WORK.
Start a new conversation. Ask a new question about the sample document and check the response.
- Keep browsing and optional cloud compute off. Run the check on the Mac, without remote phone access.
- and any other network connection.

Questions about your requirements? Write to sd@superduck.ai.

HOW IT WORKS / RECEIPTS

A receipt for the deed

An assistant that can act on files acquires a new way of being wrong: it can report work it did not do. "I have corrected the totals and saved the file," it says, and the file is as it was. The claim is fluent, confident and false, and nothing in the conversation

shows the difference.

SuperDuck does not take the machine's word for it. When the machine claims to have changed a file, SuperDuck shows the receipt: which file, and what is different in it, before and after. A claim with no receipt

behind it is a claim, and it is shown as one.

The habit this builds is the right one for working with any assistant. The question is never "What did you say you did?" It is "Show me."

FOLLOW THE DATA

Local by default, connected by choice

"Runs locally" describes the AI processing. The optional services around it behave differently.

WHAT YOU'RE DOING	WHERE THE WORK HAPPENS	WHAT TO ACCOUNT FOR
Local chat and document work	The model processes the material on the user's Mac. Conversations are stored locally.	After setup, this can work offline. Files in synced folders and system backups have their own destinations.
Installation and model download	Software and model files are downloaded to the Mac.	Setup requires network access. Download services receive the requests needed to deliver those files.
Remote iPhone access	The Mac runs the AI. A relay forwards encrypted messages between the user's devices.	The Mac must be on and connected. We cannot read the encrypted message content. This is networked use, not offline use.
Search and browsing	External sites and services receive requests when browsing is enabled.	Off by default. Consider what a search or browsing request contains before enabling it for sensitive work.
Optional LinuxToaster compute	That conversation is processed on external hardware.	Off by default. Enabling it changes where the conversation is processed.

These paths reflect the published product description, not an independent traffic audit. Policies: superduck.ai/privacy.

PRACTICAL GUIDE / THE FIRST SESSION

Your first hour with SuperDuck

Install, download a model, bring one document. A guide to a first useful hour, and what to settle before important work.

PRACTICAL GUIDE

You do not have to reorganize your digital life to try personal AI. Running an open model at home has usually meant an afternoon lost to file formats, run-times and memory limits before any work got done; SuperDuck's argument is that the afternoon should go on the work. Start with one task you understand and one question you would like help with.

Three steps to a first useful session

- 1

INSTALL SUPERDUCK.
The Mac beta is on TestFlight; check the hardware requirements first.
- 2

DOWNLOAD A MODEL.
Setup checks available memory and offers a model that fits. Allow time, and disk space, for the download.
- 3

BRING SOMETHING TO WORK ON.
A question, an unfinished draft, a short document or a subject you want to understand, with a sentence on what a good result looks like.

Try: "Here is my first draft. Which parts are unclear, and what question should I answer next?"

Give the helper somewhere to start

Put the relevant material in one folder. Tell the assistant what you are trying to achieve and what would count as help. Keep the first task small: an assumption you are unsure of, an outline, a comparison of two approaches. Treat the result as a proposal. Check the facts, ask a follow-up, decide for yourself. Before you bring in material that matters, run the offline check alongside.

Know where the files live

Keep the working folder somewhere easy to find, look through and back up. If it syncs with another service, know what that means; Apple's iCloud Drive guide explains how Desktop and Documents syncing works. The eight questions at superduck.ai are a setup checklist. Settle each one before important work.

Leave something for tomorrow

End with a short note: the goal, what you learned, the next question. It gives the next session somewhere to start. If the material belongs to an organization or is regulated, get the setup approved first.

SUPERDUCK ACADEMY / ANNOUNCEMENT

SUPERDUCK OPENS ITS WORKSHOP!

Eight pupils at a time will learn to command the new machinery of intelligence. The tools are included; the ideas must be their own.

FROM THE PUBLISHER

There is no shortage of gentlemen prepared to lecture the public on the dangers of artificial intelligence. There are fewer willing to put a machine on the table, pull up a chair, and show a person what he can do with it.

SuperDuck proposes to remedy the shortage.

The company is preparing small, practical seminars for individuals and employees of businesses, with room for eight at a time in its own offices. Companies may send a team; individuals may come on their own account. The instruction will be much the same, for intelligence is no more the property of a corporation than arithmetic is.

A workshop, not a lecture

The purpose is not to teach a clerk to write his correspondence twice as fast. It is to discover what the clerk might accomplish if he could also command the services of a researcher, programmer, illustrator, mathematician and patient tutor.

SuperDuck supplies the instruments.

There are several kinds of intelligence to choose from. Some run on the owner's Mac. Others may be called upon through Apple's services, Toasted AI from LinuxToaster, or outside providers using the owner's own keys. One may employ different models for different purposes, paying for outside computation as it is needed rather than pledging allegiance to a single supplier.

A participant will learn to bring his own documents to the workbench, examine an argument, refine a manuscript, draw a diagram, generate an illustration, or set a machine to work repeatedly until a result meets a stated test.

There are commands for improving, shortening, checking and revising work. There are scripts for jobs that must be done more than once. And there are receipts to establish whether the machine actually performed an operation, rather than merely claiming that it had.

The machine may be remarkably capable. It may also be remarkably wrong. The instructor's business is to teach the difference.

CONTINUED ON PAGE 2

A LITTLE HELP IS AT HAND

Get SuperDuck.

Put your Mac to work: one document, one question or one idea. Free beta; Apple silicon Mac required.



Scan the code or visit superduck.ai

Allow time and disk space for the initial model download. Check TestFlight for current availability and requirements.

ONE THING TO TRY TONIGHT

Ask it with the Wi-Fi off.

Install, download the model it suggests, and give it one document you know well: "Which parts are unclear?" Then turn Wi-Fi off and ask again. If the answer still arrives, the claim alongside holds. If the app stalls or asks for a connection, write to sd@superduck.ai: that is a bug report we want.

WHAT WE CAN SAY, AND ON WHAT BASIS

ARCHITECTURE

After setup, local chat and document work run with the network disconnected. Browsing and external compute are off by default. The iPhone relay carries encrypted messages between the user's devices; the Mac does the processing. Offline, SuperDuck.AI cannot disclose content it never receives.

PUBLISHED BY SUPERDUCK.AI

The Personal AI Edition · No. 01 · Free to read. This is our publication and SuperDuck is our

PRODUCT DESCRIPTION

One Mac application for Apple silicon: a model picker sized to the machine's memory, a conversation and tools for working with files. A free beta on TestFlight. Described by its maker; not independently audited.

ASPIRATION

A workshop for unfinished thinking, in which the person keeps the purpose and the judgment and the software adds capacity. Whether it feels that way is for the user to decide.

READ ONLINE

With source notes: superduck.ai · Privacy policy: superduck.ai/privacy · Letters to the editor: sd@superduck.ai

MEASURED RESULTS

None published. We have not yet released figures for speed, model quality or time to a first useful session. Treat every description of usefulness in this edition as a claim to test, beginning with the offline check above.

FOUR KINDS OF CLAIM, KEPT APART