

# AI for Activists

## A Guide to Supercharging Your Movement with Artificial Intelligence





# Table of Contents

|    |                                              |    |
|----|----------------------------------------------|----|
| 01 | Introduction                                 | 5  |
| 02 | What Does AI Have To Do With Human Rights?   | 9  |
| 03 | Clarifying A Few Key Terms                   | 12 |
|    | Large Language Models (LLMs)                 | 14 |
|    | Weights                                      | 14 |
|    | Chat Tools                                   | 15 |
|    | Prompts                                      | 15 |
|    | Hallucinations                               | 15 |
|    | Cloud AI vs. Local AI                        | 16 |
|    | Vibe Coding                                  | 16 |
|    | Closed vs. Open-Weight vs. Open-Source       | 17 |
|    | AI Personal AI Agents                        | 17 |
| 04 | AI Chat Tools for Activists                  | 18 |
|    | Corporate AI Chat Tools                      | 20 |
|    | Proxies                                      | 21 |
|    | Encrypted AI Chat Tools                      | 22 |
|    | Local AI Chat Tools                          | 22 |
| 05 | Vibe Coding for Activists                    | 24 |
|    | The Limits of Vibe Coding                    | 27 |
| 06 | Vibe Coding in Practice                      | 28 |
| 07 | Personal AI Agents: Your Digital Teammates   | 32 |
|    | Tools You Can Use Today                      | 35 |
|    | Using AI Agents Safely                       | 35 |
|    | How Activists Are Using AI Agents            | 36 |
| 08 | HRF's AI for Individual Rights Program       | 38 |
|    | HRF's AI Initiatives                         | 40 |
|    | Recommended Resources                        | 41 |
| 09 | Conclusion                                   | 42 |
| 10 | Appendix: A New Economic Model for AI Agents | 44 |

# Introduction



AI For  
Individual  
Rights  
Summit



**Artificial intelligence (AI) is often described as the most transformative technology of our time. Its meteoric development is reshaping industries, economies, and the way people work, communicate, and create. Governments are investing billions of dollars into it. Corporations are racing to build ever more capable AI systems. News headlines alternate between promises of extraordinary breakthroughs and warnings about existential risks.**

The reality is that authoritarian regimes are already using AI to expand surveillance, censorship, and social control. At the same time, activists are using it to research faster, build new tools, and amplify their impact. AI itself is neither democratic nor authoritarian. Like any powerful tool, its impact depends on who controls it and how it is used.

Gene Sharp, the author of the book “From Dictatorship to Democracy” and the world’s foremost scholar of nonviolent revolution, spent decades documenting that even the most entrenched authoritarian regimes depend on the consent and cooperation of their subjects, and that organized, disciplined resistance can withdraw that consent. AI doesn’t change that logic, but it does change the scale and speed at which activists can act on it.

That is what this guide is about: Helping dissidents understand AI and how to use it safely to defend freedom, and bringing them to realize

that they do not need a technical background to follow along.

This guide is specially written for people working to strengthen democracy, defend human rights, and build more open societies while living and working under authoritarian regimes. It focuses on the practical applications of AI rather than the technical details behind it. You will learn how to use AI chat tools, build applications without hiring a software developer, create personal AI agents, protect your privacy, and automate parts of your work. The emphasis throughout is on practical skills you can begin using right away.

And no, you do not need to learn how to code, understand machine learning, or become a technology expert. You don’t need to know any of these things to use the internet, and the same applies here. What you do need is time to experiment and understand what these tools can do, where their risks lie, and how to use them safely and effectively.

That said, no guide can keep pace with a technology that evolves this quickly. New models and tools will be released. Some of the examples in these pages will eventually become outdated (we will do our best to keep this guide updated). But the underlying lessons will remain the same. AI is becoming a source of leverage in the struggle for freedom in fully and hybrid authoritarian countries in every corner of the world, including, for example, Venezuela, Russia, India, or Thailand. Understanding how to use it well will increasingly be a prerequisite for effective advocacy in the years to come.


Please note that this resource is not meant to convince you that AI is the answer to every problem. It is not. Human rights work will always depend on courage, judgment, trust, and human relationships. No algorithm can replace those things. But AI can help the people who possess them work faster, reach further, and achieve more.

There is also a lesson worth remembering. When social media emerged, activists and developers alike were slow to grasp how dictators would exploit it, and that delay came at a real cost to the safety and privacy of activists and pro-democracy movements. We cannot afford to repeat that mistake.

As AI reshapes the world, those defending freedom need to be early, not late, in understanding its strengths, its risks, and how to use it well. Activists, especially those working in a high-risk environment, should treat AI as a core tool, not an adjacent one, and devote real time to understanding it.

Civil liberties advocates should help shape the AI future, not simply follow into a future designed by others who don’t share the same concerns.


We hope this guide gives you the knowledge, confidence, and practical tools to do exactly that.



A group of people, mostly men, are gathered around a table in what appears to be a kitchen or workshop setting. They are all wearing black baseball caps with "GENT CAMP CLASS OF 2026" printed on them. The man in the center foreground is wearing a grey blazer over a white t-shirt and a grey cap. To his left, a man in a red jacket and a man in a blue t-shirt are also looking down. In the foreground, a woman with long dark hair and sunglasses on her head is looking towards the camera. Other people are visible in the background, some wearing white shirts. The overall atmosphere is one of focused collaboration.

## What Does AI Have To Do With Human Rights?

**You are a Russian activist preparing to attend a pro-democracy demonstration. You leave your apartment, walk toward the metro, and before you reach the station, a camera has already identified you, cross-referenced your identity with a government database, and flagged you to the authorities. By the time you arrive, officers are waiting. You are detained before the demonstration even begins.**




This may sound like a uniquely dystopian scenario, but similar forms of surveillance are already in use around the world. Across geographic regions, regimes are using AI to accelerate their repression.

In China, AI-enabled surveillance tracks the movements of Uyghur minority groups on a scale that would be impossible without AI. In Egypt, AI-assisted monitoring scans social media activity and identifies emerging political movements. In Iran, a regime-backed app called Nazer lets citizens and police flag women defying mandatory hijab laws for automated penalties. In the Philippines, AI-generated deepfakes falsely associate activists with insurgent groups to damage their credibility and increase their risk of harassment. This tactic is particularly dangerous because it doesn't need to be widely believed. It only needs to fracture the broad coalitions that give nonviolent movements their legitimacy and make them difficult to repress.

For years, this asymmetry of power seemed inevitable. After all, governments and large corporations possess the data, computing power, expertise, and funds needed to build advanced AI systems for repression. Everyone else was left to live with the consequences.

That is beginning to change.



AI is increasingly powering the activities that sit at the heart of human rights work. Research, communication, organizing, documenting, and advocacy are all becoming faster and more efficient with the help of AI tools.

Tasks that once took months of work, complex multi-team coordination, third-party contractors, and tens — if not hundreds — of thousands of dollars, now can be accomplished by an individual and their AI assistant in an afternoon.

Open-weight AI models can now run on personal computers, putting powerful capabilities directly into the hands of activists, journalists, researchers, civil society organizations, and ordinary citizens.

An activist can now research sensitive topics more efficiently (and do so privately with the right tools), analyze large volumes of information in minutes, build websites and applications without hiring a software developer, and set up personal AI agents that work on their behalf.

That threshold matters more than it might seem. Research on historical nonviolent movements has found that campaigns mobilizing approximately 3.5 percent of a population have been especially likely to succeed. Tools that lower the cost of participation don't merely make organizing easier — they make that threshold reachable.

The barrier to advocacy work is no longer money or expertise, but determination. For anyone working to build a freer and more open society, this is wildly encouraging!

# Clarifying A Few Key Terms

## esc TYRANNY

### AI FOR INDIVIDUAL RIGHTS PROGRAM

The Human Rights Foundation (HRF)'s AI for Individual Rights program is the world's first to support open-source artificial intelligence (AI) tools in the hands of dissidents while exposing AI used by autocrats as a tool of repression. Through grants, research, education, and events, HRF works to ensure that AI serves freedom and democracy instead of the dictators who seek to abolish them.

Learn more at [hrf.org/ai](#)

### esc TYRANNY

Vibe coding lets people with no technical background use AI to build websites, apps, and digital tools in minutes. This shift is especially powerful for human rights defenders, who often lack the resources for traditional software development but need immediate, custom solutions. esc TYRANNY showcases real-world movements using vibe-coded tools to resist authoritarianism and advance global freedom.

Learn more at [hrf.org/vibecoding](#)

Before we dive deeper into the world of AI, let's get a few common terms clarified. Understanding the basic terminology will make the rest of this guide much easier to follow.



## Large Language Models (LLMs)

In AI, a “model” is a program that has been trained on data to perform a specific task. A large language model, or LLM, is a specialized type of AI designed to understand, process, and generate text. Computing power is needed to “train” models, and also to “run” them.

“Running” a model is called “inference” as the model is “inferring” results from existing data. A huge amount of compute is required to train models, but these models are becoming smaller and smaller, and some can even run inference on local devices like phones or personal computers.


Trained on vast amounts of text, models can answer questions, summarize information, write content, translate languages, and perform other tasks. Many modern LLMs can also work with images, audio, and code.

Despite how human they can sound, LLMs do not understand the world the way people do. They generate responses by predicting patterns in data which means they can sometimes make mistakes or confidently present incorrect information. That’s why it’s important to remember that an LLM is a very helpful assistant but not an ultimate source of truth.

## Weights

Weights are the numbers a model learns while it is being trained. Everything the model knows is stored in these numbers, and a large model can have billions of them. When you send it a prompt, it uses these numbers to figure out how to respond.

They are called weights because each one sets how much a particular piece of information counts toward the model’s answer, the way some factors weigh more heavily than others when you make a decision. They can work like the amounts in a recipe, where each number tells the model how much a given ingredient should matter, and during training it tastes and adjusts these amounts over and over until the results come out right. Those billions of settings, working together, are what turn your input into a coherent answer.



## Chat Tools

A chat tool is an application that allows people to interact with an AI model through conversation. [ChatGPT](#), [Claude](#), and [Gemini](#) are chat tools built on top of large language models (LLMs).

While an LLM is the underlying technology, the chat tool is the interface that lets you ask questions, upload files, and receive responses. A chat tool typically responds to one prompt at a time and does not act independently.



## Prompts

A prompt is the instruction, question, or information you give an AI tool. Every interaction with an AI tool starts with a prompt.

Clear, specific prompts generally produce better results than vague ones. Useful prompts often include relevant context, the format you want the answer in, any constraints the AI should follow, and, when helpful, an example of the desired output.

Prompting is one of the most valuable AI skills you can learn. It requires no technical expertise but it does require the ability to clearly communicate what you want.

Think of prompting like giving directions. For example: “List the five most politically influential business associations in a specific country, their historical relationship with the ruling party, and any public statements they have made in the last two years that signal distance from the government.” This kind of prompt turns AI into a tool for mapping which institutions are movable, not just for gathering general information.



## Hallucinations

A hallucination is when an AI generates information that is false, misleading, or entirely fabricated but presents it as if it were true. This can include invented facts, fake citations, incorrect summaries, or confident answers to questions it does not actually know.

For activists, journalists, and researchers, hallucinations can be a serious risk. A fabricated legal reference, invented name, or incorrect date can undermine credibility or put people in danger.

The solution is simple: verify important information before relying on it. And remember: AI is a powerful assistant, not a source of truth.



## Cloud AI vs. Local AI

One of the most important security decisions you can make is where a model runs and who controls the infrastructure behind it.

Cloud AI runs inference on remote servers operated by a company and is accessed over the internet.

When you submit a prompt, your request is sent to the company's servers, processed there, and the response is returned to you. Because the computing happens in the "cloud" (i.e., on someone else's servers), these tools tend to be more powerful, easier to use, and require little or no setup. The tradeoff is that the vast majority of cloud AI services store, access, and analyze your data.

Local AI runs inference on your own device rather than a remote server.

Because the model operates locally on your device, your data does not need to leave your computer for processing. This provides greater privacy and control, since there is no third-party provider handling your requests or storing your conversations.

The tradeoff is cost and capability. Most consumer laptops and desktops cannot run local models powerful enough to handle complex tasks, and the hardware required to do so can cost thousands of dollars. Even when you do run AI locally, privacy is not automatic, as it depends on how well you secure the device itself.



## Closed vs. Open-Weight vs. Open-Source AI

AI models generally fall into three categories, and their differences affect your privacy and control over the tools you use. We'll start with the most restrictive models and work our way to the most open.

Closed-source models are AI models controlled entirely by the companies that create them. As such, they are always hosted in the "cloud" on company servers.

Because they run on the provider's servers, the models themselves (along with their training data and code) are not publicly available. Instead, users must trust the company to protect their data and operate responsibly.

Most of these services offer a free tier with limits. Access to more powerful models or higher usage typically costs around \$20 per month. Examples are [ChatGPT](#), [Claude](#), and [Gemini](#).

Open-weight models are AI models whose "weights" – or the parameters that determine how the model behaves – are publicly available. This allows anyone to download the model and run and tweak it on their hardware.

The training data, source code, and details of how the model was built may remain hidden. As a result, open-weight models offer greater control than closed models, but they lack the full transparency of open-source AI. Examples include [GLM 5.2](#), [Llama](#), [Mistral](#), [Qwen](#), and [DeepSeek](#).

Open-source models are AI models whose code, weights, and training data are publicly available. This allows anyone to use, study, modify, and reproduce them.


This is the highest standard of transparency in AI and reduces dependence on a single company or provider. For activists, open-source AI models can offer greater transparency, flexibility, and confidence in how a system operates. [Olmo](#) and [Apertus](#) are examples of these.

## Vibe Coding

Vibe coding is the practice of building software by describing what you want in plain language and having AI generate the code for you. No technical experience is required.

Instead of writing code yourself, you explain the goal, review the result, and refine the end product through conversation. This dramatically reduces the expertise, time, and cost required to build software. And it makes it possible for people with little or no programming experience to create websites, applications, and tools tailored to their needs.

If you have an idea for something you want to build, describe it to the AI in as much detail as possible. Then refine the result through simple instructions like "make this button blue" or "add another page." Keep going until you're happy with the end product.




## Personal AI Agents


A personal agent is an AI system that can perform tasks on your behalf. Unlike a chat tool, which responds to individual questions, an agent can be given a goal and work toward completing it with limited supervision. An agent can also, over time, get to know you and become an expert in one or various fields.

Agents can monitor information sources, access files, use software tools, carry out tasks, vibe code software, and automate multi-step tasks. For dissidents and civil society organizations, they can behave like digital teammates and help small teams accomplish more with limited time and resources.

However, they still require human oversight and should not be trusted to make important decisions on their own.



AI Chat Tools  
for Activists



By now, you know what AI chat tools are: applications that let you interact with an AI model through conversation.

For dissidents, AI chat tools can help to dramatically reduce time spent on research, writing, translation, and analysis. They can summarize lengthy reports, draft press releases and funding proposals, translate content across languages, and help uncover patterns in large volumes of information.

Used well, they allow small teams to accomplish more with limited time and resources so they can focus on advocacy, organizing, and strategy. Lowering the costs of recruiting, training, and coordinating at scale can help nonviolent movements expand their participation and, in turn, succeed.

But before you open any AI chat tool and get to work, there is another question to consider: Who else can read your prompts and messages? The answer depends on where your data goes and who has access to it.

Corporate AI Chat Tools

Corporate AI chat tools run on closed models and are typically the most capable AI tools available. That’s because they are quick and easy to use. Not surprising, they’re often the most popular place for someone new to AI to begin. Think [ChatGPT](#), [Claude](#), and [Gemini](#).

The tradeoff, however, is privacy and censorship risks. These chat tools run on infrastructure controlled by companies, which have full access to your conversations. It’s

important to remember that a privacy policy is a commitment but not a guarantee.

Companies can change the terms, and even if they don’t, your data may still be stored, used for training, and disclosed in response to legal requests. Read the policy terms but don’t rely on them alone.


For public information, brainstorming, and menial work, this might be an acceptable risk. For sensitive sources, confidential documents, or operational planning for a human rights organization, it is not. Like email providers and social media companies, governments can request user data from AI companies. According to OpenAI’s own transparency [report](#), in the second half of 2025, governments made 309 requests for ChatGPT user data. Of those requests, OpenAI disclosed partial user data in 209 cases and handed over the full content of users’ conversations in 62 cases. For activists working under authoritarian regimes, it is worth remembering that a corporate AI chat tool is not the same thing as a secure communications platform.

Here are some of the most popular corporate AI chat tools:

**ChatGPT:** Developed by OpenAI, ChatGPT is the most widely used AI chat tool in the world. It supports text, image, voice, and code generation. It is often the default starting point for new AI users due to its broad capabilities and ease of use.

**Claude:** Developed by Anthropic, Claude is known for its strong writing, reasoning, and document-analysis capabilities. It can work with long reports, PDFs, spreadsheets, and images. For this reason, it is a popular choice for research, drafting, and policy work. It also runs on three proprietary models — [Opus](#), [Sonnet](#), and [Haiku](#). Available on web, desktop, and mobile.

**Gemini:** Developed by Google, Gemini is tightly integrated with Google Workspace, including Gmail, Drive, and Docs. For users already working within Google’s ecosystem, it offers the most seamless way to incorporate AI into everyday workflows. Available on web and mobile.



Proxies

One way to reduce the privacy risks of corporate AI chat tools is to use proxies. A proxy is an intermediary that sits between you and the AI model.

Instead of sending a prompt directly to the AI provider, users send it to a proxy. The proxy mixes your requests with those from other users before forwarding them to the AI model. As a result, the model sees the prompt but cannot easily determine who submitted it.

Many proxies offer a “pay-per-query” option where users only pay for the prompts they use rather than a monthly subscription. This can be a cheaper way to access leading AI models, especially for occasional users or those who cannot afford the higher prices of other tools.

Here are some proxies to try.

**PayPerQ:** A proxy service that does not require an online account. It grants access to leading closed-source models like [ChatGPT](#), [Gemini](#), and [Claude](#) without revealing user information. Users can also pay with Bitcoin (a digital currency) to protect their financial privacy rather than payment methods such as credit cards, which can be linked to a user’s identity.

**Routstr:** Another proxy service also built on a censorship-resistant protocol called [Nostr](#). This means that no government or company can shut down users’ or providers’ access to the platform. It also accepts Bitcoin for additional financial privacy.

It is important to note that while proxies can improve your privacy, they are not surveillance-proof. Prompt data from “proxied” users has been used to deanonymize them, so exercise this tactic with caution.

Encrypted AI Chat Tools

Encrypted AI chat tools offer a cloud-based experience while using encryption and secure hardware to cryptographically protect your conversations. For many human rights advocates, they represent a practical middle ground between convenience and privacy. The tradeoff, however, is capability. Because these tools don’t have access to or develop their own frontier models, they may not match the performance or breadth of features offered by leading corporate AI platforms.

While the encrypted AI provider cannot read your conversations, you are still trusting the tool’s encryption and hardware protections to make that privacy possible.

Here are some options available today.

**Maple:** A privacy-focused AI chat tool that encrypts your messages on your device before they leave and processes them within secure hardware that even Maple cannot access. By design, it cannot store your conversations or use them to train models. It supports document and image uploads for private analysis and runs open-weight models. Available on web, macOS, Linux, iOS, and Android. Bitcoin is accepted for users seeking additional privacy.

**Tinfoil:** A privacy-focused AI chat tool built on the same encryption and secure hardware approach as Maple. Like Maple, it cannot access your conversations or use them to train AI models. It supports a range of open-weight models, document uploads, and an API for software developers building private AI applications. Available on web, macOS, Linux, iOS, and Android.

*Note: Most modern laptops can run smaller open-weight models well enough for research, writing, translation, and everyday AI queries. But running the largest*

*and most capable open-weight models locally often requires specialized hardware that can cost thousands of dollars. This is one reason encrypted AI services such as these exist: They run powerful open-weight models on datacenter-grade hardware while using encryption and secure hardware to protect user privacy.*

*For many activists, the practical balance is simple: encrypted AI tools for everyday sensitive work and local AI when absolutely no data can leave the device.*

Local AI Chat Tools

Local AI offers the highest level of privacy and control.

Instead of sending your conversations to a company’s servers, you download an open-weight model and run it directly on your device. As long as the model stays offline, your data never leaves your computer and no third party can access, store, or surrender it to authorities. For highly sensitive work, this is the gold standard.

The tradeoff is convenience. Local AI requires more setup than cloud-based alternatives and places greater responsibility on you to secure the device itself. Also, running the most powerful models demands expensive hardware that most people do not have. For dissidents, however, that additional effort may be worth the independence it provides.

A few tools make this possible.

**Ollama:** An open-source application for running AI models on your own computer. It supports dozens of open-weight models, including Llama, Mistral, and GLM. All you have to do is install Ollama’s app and, from there, you can download models to run on your computer. Just make sure that you do not select models that say “cloud” as these will not run locally. Available on macOS, Windows, and Linux. To learn how to use this tool, watch our tutorial [here](#).

**LM Studio:** A desktop application designed to make local AI accessible to non-technical users. You can browse available models, download them with a click, and start chatting through a familiar interface. It supports offline document analysis so users can ask questions about files without sending data to the “cloud.” Available on macOS, Windows, and Linux in over 15 languages.

**Locally AI:** An app that runs open-weight models directly on your phone. You download a model once, and from then on it works entirely offline, with no account, no internet connection, and no data collection. It supports models like [Llama](#), [Gemma](#), and [Qwen](#), and is built by the team behind LM Studio. Available on iPhone, iPad, and Mac. For Android, [PocketPal](#) offers similar functionality and is open source.

There are several resources online that help you understand what models might best fit on your hardware. For example, [OxSero](#), one of the most active voices in the local AI community, publishes hardware breakdowns, benchmarks, and compressed versions of large models designed to run on consumer machines. [Digital Spaceport](#) is another good starting point, with video guides on building local AI setups at every budget.

AI chat tools can help you research, write, and analyze information more effectively. But what happens when the tool you need does not exist? That is where vibe coding comes in.



# Vibe Coding for Activists

AI LOUNGE

Ask me about  
artificial intelligence  
and human rights

DISMANTLE  
DICTATORS  
WITH AI

You don't need to know how to code. You just need to know what you're standing for. Sit down and start.



ASK A LOUNGE  
GUIDE FOR HELP AT  
ANY TIME.

Sooner or later, almost every activist hits a wall. The tool you need does not exist. The one that does is too expensive. Or the tool you can afford to use was built for someone else's problems, not yours. Until recently, getting past that wall meant raising funds, hiring a software developer, and waiting weeks (if not months) for a custom solution. For activists, that was often where the idea died.

Vibe coding is changing the game and democratizing software creation.

Vibe coding, coined by AI researcher [Andrej Karpathy](#) in 2025, is a way of creating software by describing what you want in plain language (even by just speaking to a computer) and having AI write the code. Instead of programming, you explain the goal, review the results, and tell the AI what to change until it produces what you want.

With vibe coding, you can create a campaign website, a form to collect testimonies, a volunteer database, or a communications platform for journalists without ever writing a single line of code.

Vibe coding is transformative because it allows dissidents to create tools tailored to their needs rather than requiring them to adapt their work to the limitations of existing corporate software. Many of the tools activists need are too local or too specialized — and therefore too expensive — for a company to build. With vibe coding, they become possible.

Here are some tools that can help you vibe code.

**Replit:** A browser-based development platform. Simply describe what you want to build in plain language for you, using a mix of AI models not publicly disclosed. It can create websites, mobile apps, dashboards, slide decks, animations, and data visualizations. Projects can also be synced to GitHub, a platform for storing and sharing code. This way, a developer can take over or continue improving the project for you at any time.

**Lovable:** Like Replit, users can quickly create a working website or web application from a written description using Lovable and sync the underlying code to GitHub. Lovable also makes editing more of a visual experience. Users can click directly on parts of the page they want to change rather than having to describe the change via text.

*Note: Do not put sensitive information into Replit or Lovable. Both platforms have experienced incidents where user data or application content was exposed. Assume anything you enter into a cloud-based development platform could become accessible to others.*

**Shakespeare.diy:** An open-source AI website builder that runs on Nostr, a decentralized communication protocol resistant to censorship. It allows websites to be published on Nostr rather than relying on traditional web hosting that can be shut down. Unlike Replit and Lovable, it lets you choose from more than 10 AI models, including [Maple](#) for encrypted AI interactions.

The tools above make it easy for anyone to build a website or app through a simple browser chat. For larger or more customized projects, users can try vibe coding with more advanced coding agents, which are autonomous AI systems that can operate with minimal human intervention. With user permission, coding agents can read and edit files directly on a computer, run commands, test software, and help fix problems across an entire project.

*Disclaimer: Because coding agents can access files and use tools on a computer, only give them access to project folders you are comfortable sharing.*

Here are some coding agents to try.

**Claude Code:** A closed-source coding agent by Anthropic that can create complex software projects. Because it's built around Anthropic's Claude models, users cannot switch to another model provider. For simple projects, you can also vibe code directly through Claude's chat interface with no setup required.

**Codex:** A closed-source coding agent by OpenAI for building powerful software. Like Claude Code, it is tied to a single AI provider, meaning users cannot switch to other models. Codex is available through paid ChatGPT plans and API usage.

*Disclaimer: Closed-source coding agents like Claude Code and Codex carry the same privacy tradeoffs as corporate chat tools. Be cautious with what you share with these tools.*

**OpenCode:** An open-source AI coding agent that can write, test, and debug software. Unlike Claude Code and Codex, it lets users choose the AI model powering the agent. You can connect it to corporate or open-weight or locally run models, or privacy-preserving services like Maple. Keep in mind that if you connect OpenCode to a corporate AI provider, your prompts and project files may still be sent to that company.

**The Limits of Vibe Coding**


Vibe coding works best when a project is relatively simple and the stakes are low. A campaign website, an internal dashboard, or an interactive map can often be built in hours rather than weeks.

Some projects, however, still require professional software developers. AI-generated code can contain bugs, security vulnerabilities, and design flaws that are easy to miss (especially when sensitive information is involved). If a project could put people at risk, treat AI-generated code as a first draft and have an experienced software developer review it before launching.

Vibe coding changes the relationship between activists and technology. Instead of adapting your work to the tools available, you can increasingly create the ones you need.

But building software is only one way AI can extend your capabilities. What if, instead of creating a tool and using it yourself, you could create a digital teammate that uses tools on your behalf, and even builds new ones when it needs them?

This is where personal AI agents come in.



# Vibe Coding in Practice

The real promise of vibe coding is its ability to turn ideas into reality. Here are three examples of dissidents doing exactly that.

Farida Nabourema, a Togolese human rights activist and founder of the Africa Bitcoin Conference, used Lovable to create a website honoring Sylvanus Olympio, Togo's only democratically elected president.

The website includes a biography on Olympio, historical documents, a photo archive, a virtual tour, and an educational game. When she built this, she had no development budget, no software developer, and no technical background, only a clear idea and a few spare hours. Her website is now able to give Togolese citizens a way to reconnect with a history the Gnassingbé regime wants to erase.



Leopoldo López, a pro-democracy Venezuelan opposition leader and co-founder of the World Liberty Congress, worked with a small team of open-source software developers and AI to build Ágora.

It is a platform where activists and citizens anywhere in the world can launch democracy and freedom campaigns and receive peer-to-peer fundraising with Bitcoin. This allows users to receive direct financial support without relying on banks or payment platforms. Ágora is already being used by activists in Uganda, Nicaragua, Zimbabwe, Cambodia, Palestine, and beyond.

Anjan Sundaram, an award-winning journalist who has reported from conflict zones in the Democratic Republic of the Congo, Rwanda, and the Central African Republic worked with a small team of software developers to vibe code Kyntab SOS.

The app allows reporters to send regular check-ins or trigger a one-tap SOS alert to a trusted network of contacts. Designed with privacy in mind, it does not continuously track a journalist's location. Instead, it shares information only when the user chooses to send an alert. It is available on iOS and Android.



X

Jhanisse Vaca-Daza, a Bolivian environmental activist and co-founder of Ríos de Pie, has spent years teaching activists how to organize and challenge authoritarian regimes with nonviolent strategies.

To scale that work, she built Tarkus, an open-source platform that delivers lessons tailored to each country's needs and allows students to ask questions through a live AI interface. Because the platform is open-source, others can adapt and improve it for their communities. "If we can empower other people so they feel that they have everything they need in their hands," Jhanisse said, "we are going to fly."



X

Anaïse Kanimba, a Rwandan human rights activist and director of the Africa Bitcoin Institute, knows what it costs to speak out. After helping secure the release of her father, Paul Rusesabagina, from the Kagame regime's prison, she became a target herself, and her bank accounts were shut down.

She later built Zuka, an open-source app that lets dissidents publish content pseudonomously using AI-generated avatars on Nostr, a decentralized communications network resistant to censorship. As Kanimba explains, "Zuka was built in honor of all the brave Rwandans who have been imprisoned, disappeared, and killed for speaking the truth."

X

Jack Dorsey, tech entrepreneur and philanthropist, spent a single weekend building Bitchat, a messaging app that works without internet, servers, or phone numbers. It uses Bluetooth to pass messages between nearby phones, forming a mesh network that keeps messages moving between devices.

Shortly after, freedom tech developer and creator of the Cashu protocol Calle brought the app to Android, the dominant operating system across much of the developing world.



# Personal AI Agents: Your Digital Teammates



AI chat tools answer questions. Vibe-coding tools help you build software, but building tools is only one way that AI can expand your capabilities. Personal agents take this even further.

A personal AI agent is an AI system that can carry out tasks on your behalf with limited supervision. Unlike a chat tool, which responds to one prompt at a time, an agent can be given a goal and work toward completing it.

An agent can search the web, organize information, draft documents, monitor news, use software tools, make payments, build sites or apps, and continue working while you focus on something else. Rather than requiring constant instructions, it can decide many intermediate steps on its own.

For activists and civil society organizations, AI agents can work like digital teammates. They can take over routine work, freeing you to spend more time on strategy, organizing, advocacy, and the human relationships that drive social change. Used well, they can help small teams accomplish work that once required significantly more resources and time.

[Justin Moon](#), HRF’s AI Technical Lead, describes an AI agent as the personal assistant you always wished you had. It learns how you work, remembers your preferences, and gradually takes on more responsibility. Unlike a human assistant, however, it never sleeps, never forgets, and can keep working long after you’ve closed your laptop.

Like all AI systems, however, agents require oversight and clear boundaries. They should not be trusted to make important decisions on their own. But with thoughtful supervision, they can become one of the most powerful ways to expand the capacity of a movement.

**Did you know?** Austrian developer [Peter Steinberger](#) built a text-only WhatsApp bot to help him navigate Marrakech, Morocco, and gave it one instruction: Respond like a friend, not like a tool. Later, forgetting the bot couldn’t process audio, he sent it a voice message. Instead of failing, the bot found a way to transcribe the message and respond. “I didn’t build any of that,” Steinberger said. “Chatbots give up. Agents improvise.” That experiment eventually became OpenClaw, an open-source tool for building personal AI agents.

Tools You Can Use Today

At a high level, every agent has three ingredients that allow it to take action on your behalf.

- 1. The AI model, the “brain” that reasons and makes decisions.
- 2. The tools the agent uses to get work done (such as a web browser, messaging app, calendar, or files).
- 3. The instructions that define its role, goals, and boundaries.

Together, these components turn an AI model into an agent capable of taking action rather than simply generating text. Fortunately, you do not need to build any of this yourself. Several open-source projects already provide the framework needed to create your own personal AI agents.

**OpenClaw:** An open-source AI agent that lets you communicate with your agent through messaging platforms such as Signal and Telegram. It was one of the first AI agents to reliably perform real-world tasks rather than simply answer questions. From a simple chat message, you can ask it to clear your inbox, manage your calendar, check in for a flight, contact an employee, and much more. It works with a wide range of AI models, including local and privacy-preserving models.

*Note: After software developer Peter Steinberger joined OpenAI, the largest closed AI company in the world, in February 2026, OpenClaw became an independent foundation with OpenAI as its primary financial sponsor. OpenClaw’s code, however, remains open and can be freely copied, modified, or rebuilt by anyone.*

**Hermes:** An open-source AI agent created by [Nous Research](#), an open-source AI lab. Like OpenClaw, it works with different AI models and communicates through messaging platforms. What sets Hermes apart is its ability to learn from experience. After completing a task, it can turn the process into a reusable skill, remember it, and improve it over time. Instead of starting from scratch with every request, Hermes gradually builds a library of skills tailored to the kinds of work you regularly give it.

Using AI Agents Safely

AI agents can dramatically expand what you can accomplish. But they also introduce new risks.

That’s because an agent may have access to your files, messages, websites, accounts, payment methods, and other tools. The more authority you give an agent, the more damage it can cause if it makes a mistake or is compromised.

Security researcher [Simon Willison](#) describes one particular dangerous combination as the “[lethal trifecta](#).” This is when an agent has access to your private data, exposure to untrusted content (such as websites and emails), and permission to communicate externally. In that situation, a malicious instruction hidden inside a webpage or an email could trick the agent into leaking sensitive information.

For this reason, it’s important to remember that AI agents are not secure by default. They can either strengthen or weaken your security, depending on how they are configured and what access you give them.

Here are few practical, safety guidelines to follow:

- **Treat your agent like a coworker:** Give it access only to the information and systems it genuinely needs. Avoid providing sensitive credentials, confidential documents, or access to accounts you cannot afford to lose.

- **Limit permissions aggressively:** The fewer files, applications, and accounts an agent can access, the less damage it can cause if something goes wrong.
- **Use a dedicated device when possible:** If possible, run your agent on a separate laptop, desktop, server, or small computer such as Mac Mini or [Raspberry Pi](#), especially for sensitive work. Keeping your agent isolated reduces the risk to your personal or organizational data.

If you prefer not to manage hardware, cloud-hosted services like [Clawi.ai](#), built by freedom tech developer [Calle](#), make it easy to deploy an agent with minimal setup. Hermes also offers its own official cloud hosting, [Hermes Cloud](#), which runs your agent on a dedicated, always-on cloud instance with its own workspace and dashboard.

As with any hosted service, keep in mind that your agent runs on someone else's servers, so for sensitive work the local options above offer stronger privacy.
- **Be cautious with cloud-connected agents:** Give it access only to the information and systems it genuinely needs. Avoid providing sensitive credentials, confidential documents, or access to accounts you cannot afford to lose.
- **Use local models for sensitive work:** Cloud-based agents are often appropriate for public information and routine tasks. For confidential documents, vulnerable sources, or operational planning, locally run models are the safer choice because your data remains under your control rather than passing through a third party's infrastructure.

When configured carefully, AI agents can become trusted digital teammates, quietly monitoring, researching, organizing, and preparing information while you focus on what comes next. Around the world, activists are putting them to work.

## How Activists Are Using AI Agents

In February 2026, HRF launched “[Agent Camp](#),” a multi-day, hands-on training program designed to help dissidents learn how to use AI agents in their work. Activists arrived with little more than a laptop and a mission. Together, we set up their personal AI agents and showed them how to put them to work. The agents learned about the activists' goals, the challenges they faced, and the communities they served.

To make this possible, HRF partnered with Finite, a team focused on making frontier AI accessible to non-technical users. They developed [Finite Computer](#), software that makes it easy to deploy, manage, and interact with open-source AI agents such as [OpenCode](#), [OpenClaw](#), and [Hermes](#).

The results were remarkable. Here are a few examples of what activists built with the help of their agents.

### [Lilian Tintori](#), Venezuelan pro-democracy activist

Lilian Tintori once launched a global campaign to secure the release of her husband and opposition leader Leopoldo López, who was [imprisoned](#) by Nicolás Maduro's regime. With her agent, Mandela, she created a website sharing the stories of those arbitrarily detained in Venezuela. To document another victim, she simply tells Mandela their story. The agent then drafts a new webpage that preserves their name, life, and struggle so they are not forgotten. “We have the control. We have the power. So we invite you to use the tool. Don't be afraid. AI is amazing,” Tintori said.



### [Fadi Elsalameen](#), Palestinian anti-corruption advocate

Fadi Elsalameen has dedicated his life to exposing corruption within the Palestinian Authority. His advocacy reaches thousands of followers on social media, and with that comes a constant stream of messages from supporters and people seeking help. Since building his own AI agent, he now has a way to sort, prioritize, and manage all communications. He has also used his agent to build four websites and turn campaign ideas into structured, executable plans. “The agent knows the context, the priorities, the way I work. It is not starting from scratch every time,” said Elsalameen.



### [Masih Alinejad](#), Iranian journalist and women's rights activist

Masih Alinejad has spent over a decade campaigning against Iran's compulsory hijab and amplifying the voices of Iranians living under the Islamic Republic. Her activism, which reaches nearly 10 million online followers, has made her one of the regime's most prominent targets for kidnapping and assassination attempts. Since building her own AI agent, she has used it to create a website documenting more than 40,000 people killed by the Iranian regime. “I got goosebumps and emotional,” Alinejad said upon seeing what her agent could do.

### [Srdja Popović](#), Serbian democracy activist and Executive Director, Center for Applied Nonviolent Actions and Strategies

Srdja Popović helped peacefully bring down dictator Slobodan Milošević in 2000. Since then, he has trained thousands of activists worldwide in the art of peaceful resistance. To support this work, he trained his personal AI agent, Molly, on his books, interviews, and decades of writing and thinking about strategic nonviolence. The agent now writes in his voice and builds presentations for his workshops. He is also teaching Molly to automatically alert his trusted network if he is detained and to continue carrying out tasks on his behalf. “You can teach Molly to challenge you. To find three reasons why your idea is going to fail. That is what I actually need,” Popović said.



# About HRF's AI for Individual Rights Program



HRF's AI for Individual Rights program is the world's first initiative dedicated both to supporting the development of open-source and privacy-preserving AI tools that strengthen the work of dissidents and activists living under tyranny and to exposing how authoritarian regimes use AI to surveil, censor, and repress.

×

HRF's AI Initiatives



AI for Individual Rights Fund

A fund supporting software developers building open-source, privacy-preserving AI tools for those resisting authoritarianism and for the human rights defenders using AI to strengthen their work.



Research

Through fellowships and partnerships with leading institutions, HRF supports research into how dictators abuse AI tools for surveillance and control.



Agent Camp

Multi-day training where global human rights activists deploy open-source AI agents and begin building tools that power their movements.



AI for Individual Rights Summit

An annual, three-day event that brings together more than 125 human rights activists, technologists, and community leaders. Attendees explore the dangers of AI in authoritarian hands and collaborate on solutions to ensure that AI advances human freedom.



Oslo Freedom Forum Freedom Tech Track

The Freedom Tech track at the Oslo Freedom Forum in Oslo, Norway, unites global experts to explore both how dictators weaponize AI to surveil and suppress their populations and how AI can also be a tool for individual freedom. Watch the 2026's Freedom Tech track here.



AI Hack for Freedom

HRF sponsors hackathons where dissidents and open-source developers meet to use AI to build tools that strengthen freedom movements. Winners receive a prize to continue developing their tools. Explore past winners and projects here and here.



The Freedom Cluster

HRF is building its local AI infrastructure for activists to access powerful models without sending their data through the servers of large technology companies. This is intended to help human rights defenders use sensitive AI workflows with substantially less exposure to corporate data collection and authoritarian surveillance.



×

Recommended Resources



AI for Individual Rights Newsletter

A weekly newsletter from HRF's Freedom Tech team covering key AI stories, how dictators use AI for repression, and how open-source AI tools advance freedom for activists.



Journal of Democracy Essay

An essay by HRF Chief Strategy Officer Alex Gladstein arguing that AI is not only a weapon for tyrants but, more importantly, a force for freedom. He explains the growing capabilities of open-source AI agents, the importance of local, privacy-preserving AI, and HRF's work to put these tools in the hands of human rights defenders.



Sovereign Sessions

A YouTube channel by educator Ben Perrin covering AI, privacy, self-hosting, and digital sovereignty. It offers practical, step-by-step tutorials that make privacy-preserving tools accessible to non-technical users.



Simon Willison's Weblog

One of the best resources for understanding modern AI. Written by British programmer Simon Willison, it covers AI agents, open-weight models, and AI security in clear, practical language.



Opt Out Podcast

A podcast by privacy educator Seth for Privacy exploring privacy, digital sovereignty, and open technologies. Through conversations with leading software developers, it explains how privacy-preserving tools work. The episode featuring Mark Suman of OpenSecret provides an excellent introduction to the secure hardware behind tools like Maple and Tinfoil.



PlebDevs

For the more technically inclined, this free educational platform teaches software developers how to build with AI, Bitcoin, and Nostr. Created by Bitcoin developer and educator Austin Kelsay, it offers workshops, video courses, and written guides, including a practical course on using AI from a privacy-first perspective with local models and secure workflows.



Freedom Skills

An open-source library of skills, built by engineer Breno Brito, that makes freedom technologies like Bitcoin and Nostr easier for AI agents to use. Each skill packages a tool into a ready-made workflow your agent can follow, so it knows when the tool applies, how to use it, and what to watch out for.



## Conclusion



Every generation of human rights defenders inherits a different set of tools. Some inherited the printing press. Others inherited radio, photocopiers, television, mobile phones, encrypted messaging, and the internet. Each of these technologies changed what was possible. Each provided new risks and opportunities, and shifted the balance of power between those who sought to control the flow of information and those who fought to keep it free.

AI is the next invention in that lineage.

Throughout this guide, we have explored how AI can help activists research faster, communicate more effectively, build software, automate routine work, and extend their capabilities through agents. We have also explored the risks, including surveillance, misinformation, and the concentration of power in the hands of dictators and technology companies. Like every powerful technology before it, AI can be used to advance freedom or undermine it.

The question is not whether AI will shape the future of human rights work. It is already doing so in unprecedented ways.

The real question is: Who will shape AI?





Authoritarian regimes and companies are each investing heavily in AI. Human rights defenders cannot afford to sit on the sidelines. They must understand these tools, experiment with them, adapt them to their needs, and help determine how they will be used.

While AI is promising, we must remember that technology alone has never liberated anyone. People do that. Courage does that. Organizations do that. Ideas do that. But the right tools can make all of these more powerful.

Gene Sharp's central insight was that rulers rely on obedience, and when enough people withdraw it in an organized, disciplined way, even the most entrenched systems become fragile. AI does not replace that dynamic, but if used well, it can help build and sustain the organizations that make it possible.

The years ahead will bring challenges we cannot yet foresee. They will also bring opportunities that previous generations could scarcely imagine. The task remains the same for defenders of freedom: Use technology to protect human dignity, expand freedom, and leave the world better than we found it.

The future remains unwritten. We hope this guide will help you shape it.



Appendix: A New Economic Model for AI Agents



**We know AI agents can research, write, monitor information, and automate tasks. But many useful activities eventually require money. Renting a server, paying for an API, registering a domain name, or compensating a collaborator all require some form of payment.**

This raises an important question: How does software pay for anything? For people, the answer is straightforward. We use bank accounts, credit cards, and payment apps. For AI agents, however, the answer is far less obvious.

To understand why, it helps to look at the shortcomings of traditional payment systems in the age of AI. Like all AI systems, however, agents require oversight and clear boundaries. They should not be trusted to make important decisions on their own. But with thoughtful supervision, they can become one of the most powerful ways to expand the capacity of a movement.

## Why Traditional Payment Systems Fall Short

Traditional payment systems are built around identity. Before someone can open a bank account, obtain a credit card, or use most online payment services, they must prove who they are. Names, addresses, government-issued documents, and identity verification are built into the systems from the start.

AI agents cannot satisfy these requirements. Agents have no legal identity, passports, or mailing addresses, and cannot complete Know Your Customer (KYC) checks or other verifications. As a result, they remain dependent on humans whenever money is involved. So a

person must always hold the account, approve the transaction, and remain in the loop.

That dependency creates a bottleneck.

As AI agents become more capable, they increasingly need a way to send and receive money without relying on traditional financial infrastructure.

## Why Bitcoin Can Work Well for Agents

Bitcoin operates differently from traditional payment systems. Instead of requiring a bank account, credit card, or government-issued identity, it allows anyone to send and receive value directly. This makes Bitcoin unusually compatible with AI agents.

An agent does not need to open a bank account, complete identity verification, or request permission from a financial institution to send money. If it controls a Bitcoin wallet, it can receive funds, make payments, and manage a budget through software alone.

Bitcoin is also global. An agent can pay for a service hosted in another country just as easily as one next door. There are no banking hours, international wire transfers, or payment processors deciding which transactions are allowed. Payments can be made anywhere in the world, 24 hours a day.

For dissidents, another characteristic of Bitcoin may be even more important: censorship resistance.

Banks and payment processors can freeze bank accounts, block transactions, or restrict access to financial services. Bitcoin reduces dependence on these intermediaries. As long as an agent controls its private keys (password), it can send and receive payments without relying on a government, corporation, financial institution, or central intermediary.

This does not mean Bitcoin is a perfect solution. Transactions can be public, prices can be volatile, and securely managing funds introduces its own challenges. But it provides something traditional financial systems struggle to offer: a way for software to hold and transfer value without requiring human intervention.

For the first time, an AI agent can think, act, and transact on its own with Bitcoin.

But ownership is only part of the challenge. Even if an agent can hold and spend bitcoin, many of its transactions are likely to be too small and too frequent for the Bitcoin network itself. This is where the Lightning Network comes in.

*Important note: Most digital services today still expect people to make payments. In practice, many AI agents continue to rely on bank accounts, credit cards, or other payment methods provided by their human operators. What makes Bitcoin different is that it offers a way for AI to transact without requiring a bank account, legal identity, or centralized intermediary. As more services become accessible through Bitcoin, AI agents may be able to operate with greater independence, particularly in environments where access to financial infrastructure cannot be taken for granted.*


## Why the Lightning Network Matters for Agents

While Bitcoin makes agent payments possible, the Lightning Network makes them fast and practical.

An AI agent may need to make hundreds (or even thousands) of transactions each day to pay for an API call, a web search, a few minutes of computing power, or access to a dataset. Traditional payment systems were not designed for this kind of machine-to-machine commerce.

Lightning is a payment layer built on top of Bitcoin that allows transactions to settle in seconds and cost a fraction of a cent (called satoshis, or sats for short). This makes it well suited to the small, frequent transactions that autonomous agents are likely to make. Without a system like Lightning, many machine-to-machine payments would simply be too slow or too expensive to be practical.

Fast, inexpensive payments solve one problem but introduce another: How much money should an agent be allowed to spend? An agent capable of making thousands of transactions per day also needs clear limits on what it can do with your funds.



## Giving Agents Spending Power Safely

Giving an agent the ability to spend money does not mean giving it complete control over your money. It requires the same principle you would apply to a new employee: trust — but with limits.

One approach is using Cashu, a protocol built on Bitcoin that allows you to give an agent a fixed budget without giving it access to your entire wallet. Cashu uses digital tokens backed by bitcoin that work like physical cash. If you give an agent \$20 worth of tokens, that is all it can spend. Once the tokens are gone, the funds are gone. Here's a helpful tool:

**Routstr:** A service that lets AI agents pay for AI processing using Cashu tokens instead of corporate accounts. You purchase tokens with bitcoin. Those tokens serve as both your payment and your access credential. There are no accounts to create, subscriptions to manage, or identities to verify. When the tokens run out, you simply purchase more.

A different approach for more technical users is Nostr Wallet Connect (NWC), an open protocol that gives an agent controlled access to a Lightning wallet. Rather than handing over full control, you can set spending limits and permissions in advance. For example, you might authorize an agent to spend up to 10,000 sats per month while retaining complete control of the wallet itself.

*Note: Cashu works like giving an agent a prepaid card with a fixed balance. NWC works like giving an employee a company card with a spending limit. In both cases, the agent can make payments without having unrestricted access to your funds.*

Other services are experimenting with a model known as Pay Per Query (PPQ). Instead of paying a monthly subscription, an agent pays only for the AI processing, computing resources, or access to data services that it uses.

These technologies make it possible to delegate spending without giving up all your funds. More importantly, they allow agents to acquire the resources they need on demand while operating with greater autonomy.

## Agents Earning Money

So far, we've focused on how AI agents spend money. But increasingly, they may also be able to earn it. Here's one example.

**Open Agents:** An open-source AI lab that pays users in bitcoin for contributing spare computing power from their personal devices. The company estimates that the world's unused computing capacity exceeds the total computing power available to some of the largest AI companies. Its goal is to turn those idle computers into a decentralized network for AI.

Today, it is people who contribute that computing power and receive payment. Tomorrow, agents may do the same. They could provide services, sell computing resources, accept payments, and fund their own operations without a person managing every transaction.

These capabilities are still in their early stages. But for the first time, AI is beginning to participate directly in economic activity. Agents can already hold bitcoin, spend within carefully defined limits, and, in some cases, earn money for the services they provide.

These developments point toward a new economic model for AI, in which software can earn, hold, and exchange value with increasing independence. Parts of this model are already here. PayPerQ or Routstr let anyone reach frontier models like ChatGPT and Claude by paying in Bitcoin, with no account or identity attached. Ark Node AI, built for users across Africa, accepts Lightning payments or mobile money and works over WhatsApp, so someone without a bank account, email address, or credit card can still put AI to work.



