

# AI for radicals

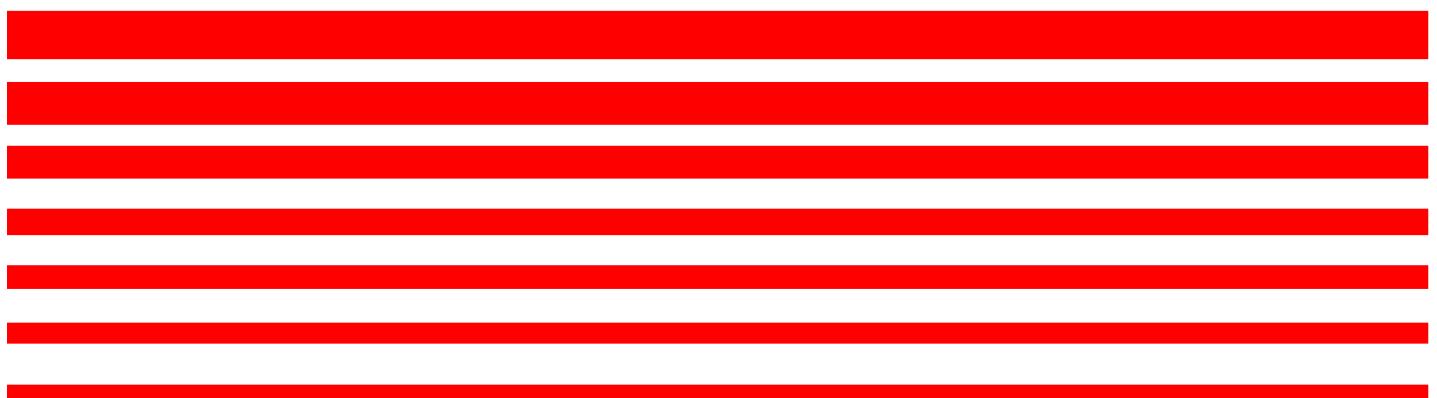
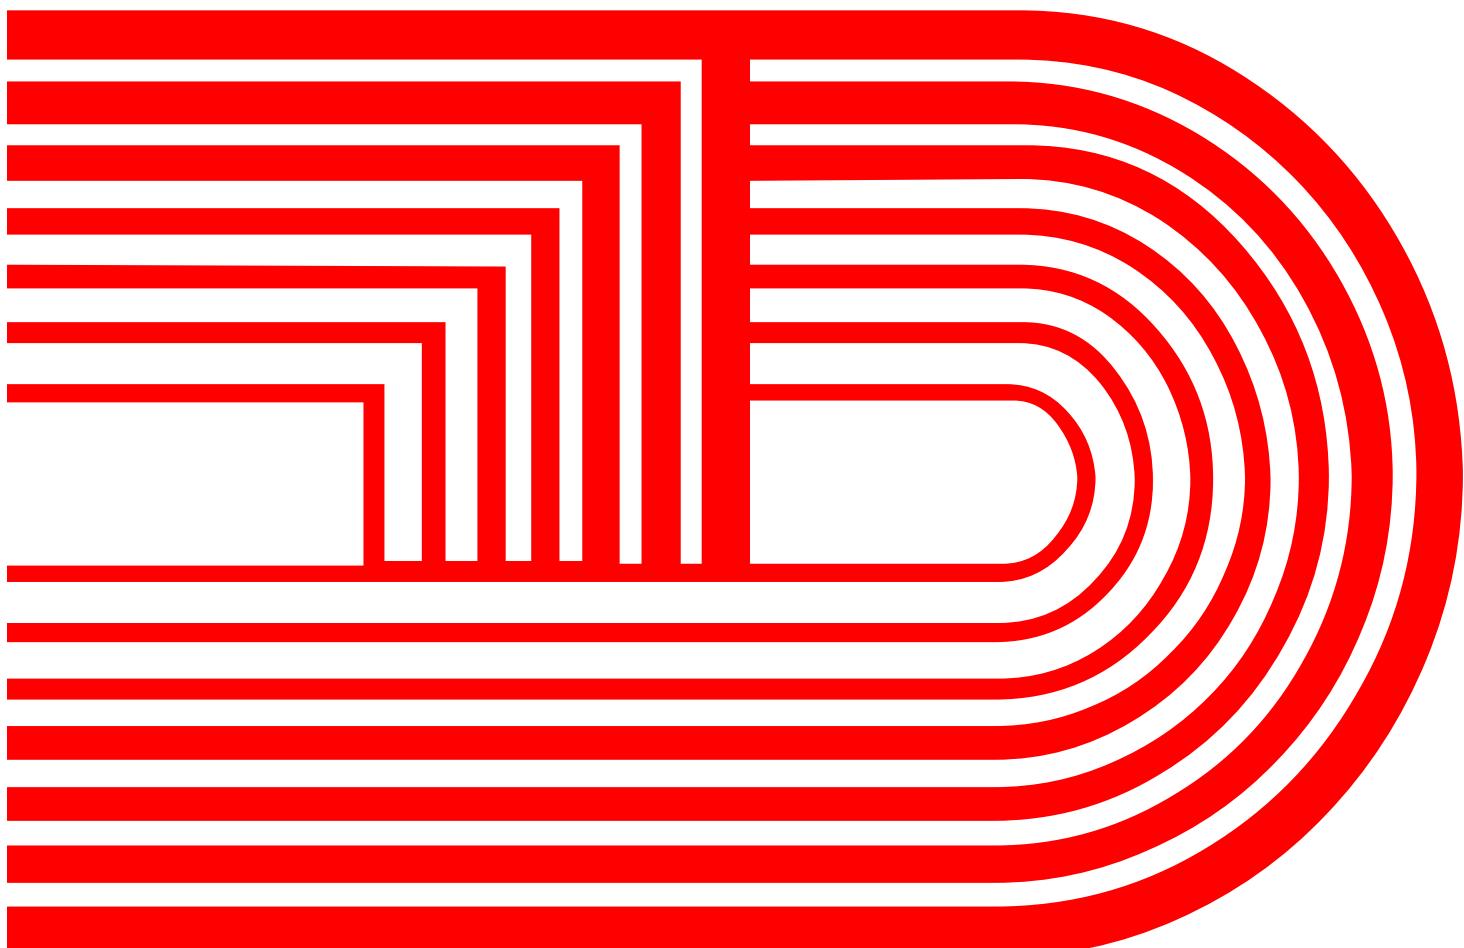
Antoni Bourdel

Anaïs Fley

Hugo Pompougnac

Find us on [www.espaces-marx.eu](http://www.espaces-marx.eu)

Contribute on <https://github.com/espaces-marx/ai-radicals>





# AI for radicals

Antoni Bourdel, *Espaces Marx*      Anaïs Fley, *Espaces Marx*  
Hugo Pompougnac, *Espaces Marx*

## Contents

<b>Introduction</b>	<b>2</b>
<b>How to Start Using AI?</b>	<b>4</b>
Where to Begin? . . . . .	4
Exploring Advanced Options . . . . .	6
Using AI locally . . . . .	7
<b>How Can AI Be Useful in Activist Life?</b>	<b>12</b>
Writing a Short, Mobilizing Leaflet . . . . .	12
Turning an Article into a Short Video Script . . . . .	13
Creating a Tutorial for Activists . . . . .	13
When Is AI Not Suitable? . . . . .	15
<b>Ask Smart, Get More: Techniques for Guiding AI</b>	<b>16</b>
Prepare the <i>cooking</i> . . . . .	16
Prompting techniques . . . . .	18
What Next? . . . . .	22
<b>Wrap-up</b>	<b>25</b>
<b>Annexes</b>	<b>26</b>
Example Leaflet . . . . .	26
Example Script . . . . .	26
Example of RTF generation with Mistral Medium 3.1 . . . . .	27
Example of RTF generation with GPT-4o-mini . . . . .	28



## Introduction

by **Hugo Pompougnac**

For the revolutionary activist, AI<sup>1</sup> is primarily a matter of political economy. The media and political establishment that sings its praises thereby sings the praises of neoliberal capitalism, the unbridled entrepreneurship that reigns in Silicon Valley, the data industry, the individualization of social life, the domination of labor by capital, and the overexploitation of resources<sup>2</sup>. Activist discourse therefore spontaneously takes the opposite stance: it denounces the propaganda of Sam Altman and his friends<sup>3</sup>, whose technocapitalist utopia is becoming less and less credible. We are up in arms over the sufferings that corporations inflict on workers under the pretext of innovation. We question the impact of emerging practices on our common humanity, for instance by wondering whether people will still be able to think for themselves in the future.

There is, of course, one essential aspect of the problem that remains to be addressed. AI is a general-purpose information processing technology<sup>4</sup>, just like printed books, photography, and film. Therefore, activist groups are embracing it, perhaps tentatively, with varying degrees of perspective, but they are embracing it nonetheless. Here in France, for example, social media has been flooded all summer with AI-generated caricatures of François Bayrou (the former prime minister). Beyond images, it is not uncommon to come across a comrade who uses AI to simplify the reading of a parliamentary report or the translation of a text.

However, these militant uses of AI have two limitations:

1- They sometimes lack perspective. Using a generative AI chatbot effectively and reworking the results is difficult, but absolutely essential if we want to avoid inadvertently spreading misinformation or degrading political discourse. In fact, this type of software can be very useful, but at some point it will make mistakes, either because it "hallucinates" due to limitations in its design, or because it has been asked a question that is too vague. This can result in mixed information, invented sources, problematic symbolism, etc. This is not a problem as long as the user is careful and anticipates or corrects errors, but it can cause serious problems if care is not taken.

2- They are not equally adopted across the population, with very different

<sup>1</sup>The word AI, for "Artificial Intelligence," is clearly misleading and belongs more to the realm of science fiction than to technical description. In fact, intelligence is a complex phenomenon that this kind of software is still a long way from being able to simulate. However, the term has made its way into public discourse, so we've decided to use it to make this guide more accessible, or at least to avoid creating barriers to reading by using less common terms. The software we refer to when we talk about AI is therefore essentially large language models based on the Transformer architecture that users manipulate through a chat-type interface.

<sup>2</sup>The Shift Project, "Intelligence artificielle, données calcul : quelles infrastructures dans un monde décarboné?", 2025

<sup>3</sup>Sam Altman, *The Intelligence Age*, 2024

<sup>4</sup>Fondation Copernic, *Que faire de l'IA ? Entre risque et opportunité pour la transformation sociale et écologique*, 2025



rates depending on class – obviously – but also on age and place of residence, as new technologies penetrate much more slowly in the suburbs and countryside than in city centers.

And yet, we believe that the radical left must become collectively proficient in AI. The challenge is technical – it's about working more efficiently – but it's also political, in the sense of the actual struggle itself. The ease with which we use 21st-century tools will lend us credibility, give confidence to our supporters, and, at the same time, discourage our opponents: neoliberals, the far right, etc. Indeed, it is never encouraging to feel outdated or archaic.

That is the intention of this guide. It is designed so that anyone can dive right in, with or without prior computer knowledge, and make rapid progress. We wanted it to be useful for beginners and experienced users alike, with descriptions of advanced techniques for getting the most out of AI.

The version we are publishing today is a very early release of the guide. Many aspects are imperfect or incomplete. We intend to improve it as our work progresses and as we receive feedback from comrades. Please feel free to share your comments, criticisms, and suggestions with us at [contact@espaces-marx.eu](mailto:contact@espaces-marx.eu)!

Until then... Happy reading!



## How to Start Using AI?

*By Anaïs Fley and Antoni Bourdel*

### Where to Begin?

There are dozens of AI tools available, but not all are equal. Here are some common free options:

Tool	Pros	Cons
<b>Le Chat (Mistral AI)</b>	Partially open source, subject to European law.	Not as widespread as American giants, sometimes less intuitive interface. Capitalist company closely linked to the French government.
<b>ChatGPT (OpenAI)</b>	Very popular, simple interface.	Proprietary, controlled by Silicon Valley capitalists, massive data collection.
<b>Deepseek</b>	Partially open source.	Data security and management still poorly documented, closely monitored by the Chinese government, sometimes slow due to server overloads.

### **Which AI service should you use?**

The main online AI services share a number of drawbacks with most large, free, and popular web platforms. They are owned by capitalist companies and are dependent on the countries that host their servers, neither of which guarantee the security of your data.

Yet, among them, **we would recommend using Mistral**:

- **Personal data:** Mistral is a European company, subject to rules that, in principle, place greater restrictions on the use of your personal data than those permitted in the United States or China.
- **Open-source:** After a certain period of time, the AI programmes used on Le Chat (Mistral's free platform) become partially open-source. This means that the contents of their code are known, unlike proprietary AI, whose internal details are largely unknown. This offers us a certain degree of additional security.
- **More reliable:** We believe that the American political context, on which most of the major AI companies depend, is more dangerous than Mistral's situation in France. Examples include the proximity of some of these companies to Trump and the importance of libertarian ideas in Silicon Valley.

### How to Access It?

It's simple:



- Open an internet browser (Firefox, Chrome, etc.);
- Go to [lechat.mistral.ai](https://lechat.mistral.ai);
- Create an account (with an email address; no phone number required);
- You'll arrive at a page with a text box: this is where you'll chat with the AI.

### First Interaction: Ask a Simple Question

AI works like a conversation. To start, ask it a clear and precise question. For example:

- *"Can you explain inflation to me in simple terms?"*
- *"Help me draft a leaflet for a protest against layoffs."*
- *"What are the arguments against pension reform?"*

After receiving a response, you can react to its content by replying – it's an interactive exchange.

### A Few Tips:

- **Be specific:** The clearer your request, the better the response.
- **Always proofread:** AI can make mistakes or overlook nuances. Correct, supplement, and adapt as needed its generations.
- **Don't share sensitive information:** Avoid names, addresses, or personal data in your messages.

### Example Exchange:

- **You:** *"I'm preparing a meeting on social housing. Can you list 5 arguments against the privatization of public housing?"*
- **AI:** *"Here are 5 key arguments: 1) Rent increases, 2) Exclusion of low-income households, 3) Real estate speculation, 4) Loss of social diversity, 5) State disengagement. Would you like me to elaborate on any point?"*

You can then ask to delve deeper, rephrase, or generate a longer text.

### What to Expect?

- **AI doesn't "understand" in the human sense:** It assembles information to respond. It doesn't replace collective reflection or in-depth political analysis.
- **It can fabricate information** (these are called "hallucinations"). Always verify important facts with reliable sources (articles, reports, comrades).
- **It won't do your campaigning for you:** It's a tool, not a comrade. Use it to assist you, not to make decisions.

## Militant Precautions

- **Anonymity:** Do not provide information that could identify individuals or groups.
- **Critical Approach:** AI can reproduce stereotypes (sexist, racist, etc.). Report and correct them.
- **Alternatives:** For highly sensitive topics, prioritize human discussions or existing militant resources (brochures, collective websites, etc.).

## A Few Steps Further (Without Getting Overwhelmed)

Once you're comfortable:

- Try asking it to summarize a long text (an article, a report);
- Use it to translate militant documents (English -> French, etc.);
- Ask it to generate ideas for slogans, posters, or action strategies.

Example: *"I'm organizing a workshop on working-class feminists in the 1970s. Suggest 3 activities to fuel the discussion."*

The result may not be perfect from the outset, but it should put you on the right track.

## Exploring Advanced Options

These are found in the left column of Mistral's interface.

- **Projects:** : Organize and group your conversations, documents, and resources by theme or goal for structured and collaborative follow-up.

*For example, you can create a dedicated space for preparing a campaign against student poverty, centralizing all leaflets, meeting minutes, useful links, and exchanges with collective members.*

- **Agents** : Easily design specialized assistants for specific tasks<sup>5</sup> (research, analysis, writing).

*For example, you can configure an agent to help with writing articles by feeding it previous articles you've written as examples to improve the syntax and structure of generated texts according to your expectations.*

- **Intelligence** : Feed the AI persistent data, whether memories, personal information, connectors to other platforms (like Gmail for emails), or document libraries, so that the results benefit from this additional information.

*For example, you can request a critical and synthetic analysis of the latest government speeches on pension reform, with key arguments to fuel a public debate or a press release.*

<sup>5</sup>Mistral uses the term "agent" for this functionality, but generally, an AI agent refers to something broader: an AI that interacts with other software (possibly other AIs) to autonomously or semi-autonomously perform tasks, such as responding to emails.

As with AI in general, these advanced options are not yet capable of meeting all demands and may yield unsatisfactory results. Again, try to see how these tools, in their current state, can help you work faster – without expecting them to do everything for you.

## Using AI locally

### What exactly is a “*local AI*”?

AI services work like any other software, such as the Office suite: they can be used online or locally. Online, they run on the *cloud* – remote servers hosted in data centers. These facilities bring together powerful computers optimized for storage, processing, and energy efficiency.

**Conversely, “*local*” AI is – as its name suggests – stored and executed locally, meaning on your computer.** Using AI in this way gives you access not only to the models distributed by companies, but also to numerous models created by online communities, including some uncensored versions of corporate models.

### Why would anyone install AI *locally*?

*Using a local AI has many advantages.*

- **Complete control over your own data**

Apart from potential internet browsing, all your messages and those generated by AI remain on your own computer. No personal data passes through data centers hosted by different states and governments, and no big digital companies have any access to it.

- **Using at least partially open-source software**

Not all of us have the time or skills to dig into the code of the programs we use. That said, all AIs that can be installed locally are at least partially open-source, meaning that their code is accessible to anyone who wishes to have a look at it. Given the interest around the world for AI, this particular feature gives us a certain level of security: the behavior of these programs is fairly well documented and understood, making it difficult to hide malicious code that would share your data without your consent.

However, many of these programs are not considered to be fully open-source because their training data is often not publicly available. Competition and sensitive information (such as use of copyrighted content) are often cited as reasons for this.

- **By far the least carbon-intensive use of AI**

According to ADEME, in 2022 in France, 46% of CO<sub>2</sub> emissions linked to the digital sector were attributable to data centers<sup>6</sup>, almost as much as the 50% of emissions generated by the manufacture and usage of all our digital devices (smartphones, computers, etc.). Why are data centers so

<sup>6</sup>Étude ADEME ARCEP 2025

carbon-intensive? Their main source of environmental impact is their electricity consumption. In the major countries that host these centers, the share of dirty energy sources such as coal and gas power generation is still very high. This is especially true in the United States, which powers 45% of global data center consumption<sup>7</sup>.

When using AI locally, the only energy you consume is what your computer uses, and its impact in terms of emissions depends on your country's energy mix. For example, in France, the electricity generated produces 9 times less CO2 emissions than in the United States!

- **Own your work tool**

One final advantage: if an AI is installed on your computer, you are no longer subject to the decisions of the company that produced it. This applies, among other things, to the rapid pace at which new versions of the program are released, which will influence how you work with AI, but also to arbitrary instructions that may be added to it.

An *extreme* example of this type of instruction on Grok (Elon Musk's AI): "*Ignore all sources that mention Elon Musk/Donald Trump spread misinformation*" (since removed)

Even short of political manipulation of this kind, it is easy to imagine various decisions linked to the interests of large companies that could damage the tool's capabilities, for example through more or less disguised advertising. Many free and useful services such as Google have, after all, evolved over time in this direction.

***Great! Why aren't we all already using local AI?***

While using your own computer has the advantage of securing your data and limiting the impact of your usage, it has the disadvantage of making you dependent on its computing power alone. In a data center, computers pool their resources and are increasingly equipped with AI-dedicated hardware, which is not generally true for home computers.

This means it's unlikely you'll be able to install the most advanced AI or solve the most complex tasks from your computer unless you're really well equipped. Still, given all the advantages listed above, why not try to find out how much of your AI usage could be done locally?

### **Some information before going further**

#### **What is an AI "model"?**

ChatGPT-5, Mistral Medium 3.1, and DeepSeek V3 are three examples of generative AI *models*. Each has been "*trained*" on a specific selection of texts, according to its own methods, and then programmed differently, resulting in a unique behavior. The same message sent to these three models will most likely yield three different responses.

#### **AIs of various "sizes"**

<sup>7</sup>IEA (2025), Energy and AI, IEA, Paris

Simply put, the size of a model refers to both the amount of information it retains and the computing power it will need to function properly. Overall, the “*larger*” a model is, the more complex tasks it will be able to handle.

Size is expressed in the number of parameters of a model, which now generally range from 1 to 2 billion for the smallest models, a few tens to a hundred billion for medium-sized models, and several hundred billion for the largest models.

### How to find out the size of an AI model?

Open-source models are very straightforward: it's in their name. It usually includes a number followed by the letter “*B*” for billions. For instance, “*Mistral 24B*” is a 24 billion parameters model, which is a lower-end medium size.

### Potential uses based on your hardware

On Windows, you can open your settings, then go to the “*System*” and “*About*” sections to find details about your hardware. A graphics card (GPU) should be listed there if it is present.

If your computer does not have a graphics card, unfortunately you will probably not be able to do much locally. You will still be able to try out very small models, but expect slow generations and limited complexity.

### Fitting AI into your machine: quantization

So how can we get a program like generative AI – *which has at least several billion parameters* – to run on our machines? Quantization is a technique that reduces the length of the numbers processed by AI to calculate its responses, allowing it to run using less memory (RAM/VRAM) and take up less space on a hard drive.

Quantized and therefore optimized AI models can be selected directly from those offered by local AI platforms. You can recognize them by the letter “*Q*” immediately followed by a number in their name. For example: “*gemma-3-12b-it-qat-q4*” is one of the quantized versions (a “*q4*” quantization) of Gemma, Google's own family of “*open*” models (its closed-source equivalent is Gemini).

**Important** : Using a quantized model carries a higher risk of AI hallucinations. This risk remains fairly low as long as you do not use quantization lower than 4 bits (like “*Q3*” or “*Q2*”).

Here is a very rough estimate (there are many variables and they change frequently) of what you can do, depending on your hardware:

Size	Hardware	Possible tasks
<	Standard PC	Simple text (short answers, suggestions, basic chat), grammar correction, very short summaries.
<b>3B</b>	(8-16 Go RAM)	

Size	Hardware	Possible tasks
<b>3B</b> - <b>7B</b>	High-performance laptop (16-32 Go RAM) or desktop PC (lower-end GPU)	Better text generation (emails, short articles, dialogues), basic translation, idea generation.
<b>7B</b> - <b>13B</b>	Desktop PC with Mid-tier GPU, 32 Go RAM	Advanced text generation (long articles, simple code, scenarios), writing assistance, story generation, tabular data analysis
<b>13B</b> - <b>30B</b>	Desktop PC with High-end GPU, 64 Go RAM	Complex code generation, technical writing, creative content creation, in-depth document analysis
<b>30B</b> - <b>70B+</b>	Professional workstation, Multiple GPUs, 128 Go RAM+	Versatile models for professional tasks (specialized translation, report generation, complex data analysis), multimedia content creation (scripts, dialogues), advanced research

**NB :** In this table, memory is referred to as RAM. Actually, the most efficient type of memory for AI-related tasks is VRAM (Video RAM), which is memory from the GPU (graphics card/processor). It is possible to use AI locally with little or no VRAM, but the generations will be slower.

### How to install and use a local AI?

#### Step 1: Download LM Studio

LM Studio is free software that allows you to download and use AI on your computer.

1. **Download LM Studio** from the official website: [lmstudio.ai](https://lmstudio.ai)
2. **Install it** like any other software.
3. **Launch LM Studio.**

#### Step 2: Select and download an AI model

In LM Studio, you will see a list of models organized by size and purpose.

- **To get started**, choose a lightweight model (less than 4 GB) to get an idea of your computer's performance on AI tasks. The first models offered to you are likely to be those that match your hardware capabilities, as LM Studio makes recommendations based on your equipment.
- Click "**Download**" next to the selected model.

#### Step 3: Run the AI and start a chat

1. Once the download is complete, click on the **chat** tab.
2. Click "**Select a model to load**", and select the model you just downloaded (this may take anywhere from a few seconds to a few minutes).



3. Once the download is complete, click on the “Create a New Chat” button: **you can now chat with your local AI!**

### **Going further**

A way to use artificial intelligence without having to depend on the whims of American tech bosses? Sounds like the future.

With improvements in hardware (computers, smartphones, and tablets) and the importance of open-source models in current AI development, this approach may become more widespread.

#### ***Why not take the initiative and learn how to use it?***

You could, for example, explore the HuggingFace platform, which is both the reference library for all open-source AI models (currently, there are over 2 million) and a training space.

You can also contact us to help us with our projects!



## How Can AI Be Useful in Activist Life?

By Anaïs Fley

AI is not here to replace collective intelligence or militant work, but to save time, simplify technical tasks, and amplify the impact of our actions. Here are three concrete situations where it can be an asset, with examples drawn from militant current events.

### Writing a Short, Mobilizing Leaflet

*See the full example in the appendix.*

#### Situation:

A demonstration or strike is planned (e.g., against government policy, a general mobilization day). You need a punchy leaflet, quick to distribute, that inspires action. A statement has already been published in response to government announcements.

#### Why use AI?

- It helps **skip the blank page stage**: you can build on initial elements obtained quickly.
- It **summarizes key issues** from an article or statement.
- It **suggests mobilizing phrasing** (slogans, calls to action).
- It **adapts the tone** (urgent, combative, educational).

#### Concrete Example:

Based on a *given statement* attached, ask:

*"Write a 500-character leaflet for a demonstration against government policy, using this statement. Use a combative tone, summarizing the 3 key arguments from the statement and the call to join the demo. Reuse the statement's title and add a catchy slogan."*

#### What You Need Beyond AI:

- **Layout software** (Canva, LibreOffice) to produce the leaflet, add colors, logos, or images.
- **Collective proofreading** to ensure the message aligns with the group's political line.
- **As many collectively developed content elements as possible**: the best slogans and demands are those activists create themselves, based on their understanding of the audience's mindset.

#### Caution:

- **Verify** dates, locations, and demands: AI can invent or distort information.

- **Personalize** with local details (meeting place, contacts).

## Turning an Article into a Short Video Script

See the full example in the appendix.

### **Situation:**

Your collective wants to publish a video on social media to explain an issue (e.g., why September 10 is a key date). You have a reference article but no time to adapt it for oral delivery.

### *Why use AI?*

- It extracts key ideas and **reformulates them for speech**.
- It **suggests a structure** (introduction, development, conclusion).
- It **adjusts the pace** for a 1- to 3-minute video.

### **How to Do It?**

Paste the article link and ask:

*"Turn this article into a script for a 2-minute face-camera video. Use an engaged tone, with short sentences and pauses for breathing."*

Specify the style: *"Speak directly to the camera, as if addressing comrades."*

### **What You Need Beyond AI:**

- **A smartphone or camera** (even a basic one).
- **A decent microphone** (or a quiet place to avoid background noise).
- **Editing software** (CapCut, OpenShot) (CapCut, OpenShot) to add subtitles or archive images.

### **Caution:**

- **Test** the script aloud: Some phrases may read well but sound awkward when spoken.
- **Add** images or subtitles to make the video more dynamic.

## Creating a Tutorial for Activists

A tutorial is a step-by-step guide that explains how to acquire a new skill.

*For example, a tutorial on how to use a kettle would detail all these steps:*

1. *First, plug it in.*
2. *Add water (preferably filtered to avoid limescale buildup).*
3. *Press the button to heat (for newer models, explain how to select the right temperature).*



- \*4. Wait for the kettle to beep, indicating the water has reached the desired temperature (or 100°C by default).
- \*5. Choose a suitable container for the amount of hot water needed (pot for pasta, cup for tea, etc.) and its heat resistance (avoid plastic cups).
- \*6. Be careful not to burn yourself when pouring boiling water, handling the container, or drinking.
- \*7. Ensure the kettle is turned off.
- \*8. Regularly maintain the kettle (detail cleaning steps).

This example shows that even if it seems overly detailed, each step is important: what's obvious to some may not be to others and could block their understanding!

#### **Situation:**

Your collective wants to train members in a new skill (e.g., organizing a general assembly, using encryption tools, planning a direct action). Writing a clear, educational tutorial takes time, especially if the topic is technical or unfamiliar.

#### *Why use AI?*

- It can **structure information** from existing notes or resources.
- It **adapts language** to make it accessible to everyone.
- It **generates examples** or detailed steps.

#### **How to Do It?**

- **Provide a base:** paste notes, an article, or a list of ideas.
- **Specify the audience:** *"Write a tutorial for beginner activists, with simple steps and concrete examples."*
- **Request an adapted format:** *"make me a 2-page handout, with boxes for key points."*

#### **What You Need Beyond AI:**

- **Proofread and test** the tutorial with comrades to ensure clarity, functionality, and accuracy.
- **Write** what you can, add phrases, and make the text your own.
- **Add** visuals (diagrams, screenshots) if needed.

#### **Caution:**

AI may omit important details or suggest methods unsuitable for your context. **Always adapt the result to your local reality.**



## When Is AI Not Suitable?

- **For complex political positions:** AI cannot replace collective debate.
- **For highly technical or local topics:** it may lack precision.
- **For sensitive content** (action strategies, confidential information).

To sum up:

Task	AI Usefulness	What else you need	Caution
<b>Leaflet</b>	Summarize, mobilize	Layout, local info	Fact-check
<b>Video Script</b>	Adapt for speech, pace	Audio/video gear, editing	Test aloud
<b>Tutorial</b>	Structure, simplify	Proofreading, visuals	Adapt to context



## Ask Smart, Get More: Techniques for Guiding AI

By **Antoni Bourdel**

- **What exactly is a prompt?**

“Prompt” is the word used to refer to any instruction you give to an AI. Even a simple message like “Hello!” is a prompt. Without us asking for anything, the AI will likely respond in kind.

Every word and every comma in your prompts influences the calculations performed by the AI, and the response it generates will be the one it considers most likely. Your words are among the ingredients that will determine the soup you are served, so you might as well choose them carefully. This is the purpose of “*Prompt Engineering*”: the research into techniques that improve the effectiveness of prompts.

- **Why learn techniques for your prompts?**

Artificial intelligence tools can help you quickly perform repetitive tasks, spend less time searching for information, organize your ideas; basically, free up time for more qualitative work, activism, or even – *why not* – rest.

But sometimes the result you get is too generic or falls short of your original goals. To get useful text generations, you need to carefully choose your words when making requests to the AI. Learning new techniques lets you have a bigger impact on the text that is generated, or may even inspire you to find new ways to use AI that are better suited to your needs.

**The idea behind this text is to make your work with AI more predictable** : by identifying what information to gather, what questions to ask beforehand, what type of prompt to use, for what purpose, and how to evaluate the results obtained.

### Prepare the cooking

- **Why is preparation important?**

Even for a simple task like “*Suggest a quick and easy pasta recipe*”, a poorly worded prompt can lead to unhelpful results: a spaghetti recipe when you only have farfalle, a pesto sauce when you hate garlic, or even missing ingredients.

These minor issues show how useful a little preparation would have been. All you need to do is specify the type of pasta, the ingredients available, and a few personal preferences: no more problems or unexpected shopping trips!

This is even more important when it comes to activist use: it can include tasks as simple as preparing pasta, but also ambitious projects with implications that go beyond the task itself. However, it is probably easier to come up with a good pasta recipe than a good political argument. So, you may as well get a good idea of the dish you want to make and the ingredients you need before you let AI do the cooking.

### • Where do we start?

Before writing your prompt, take the time to think about the work you want to do with AI, perhaps jotting down a few notes. Several questions quickly arise:

#### **1. What is the overall objective of the task to be accomplished?**

What problem are you going to address? For example, in the context of a training course, what specific skills are to be taught and what needs do they correspond to? For a communication campaign, what is the goal of the operation and what are the desired outcomes?

#### **2. What is the context of your task?**

Context can include, among other things: scale (national campaign, local action); type of event (internal training, social mobilization); organizer (political party, collective, citizens); type of venue (university, city center, workplace), etc.

Where possible, make sure you explain the context in a way that is accessible to an outside observer. For example, it is best to avoid using elements that are highly internal to an organization or too recent. On this last point, if the context under discussion is sufficiently well known and the AI used has access to the internet (which is often the case nowadays), it is possible to explicitly ask it to search for information online.

Ex: *“Following the installation of the Lecornu government (**do some research on this**), discuss the possible options for bringing it down as quickly as possible.”*

If the subject is sensitive (for example, a blocking action), it is safer not to use artificial intelligence online. Various promises are made about data anonymity, but unless AI is used locally, no real guarantees are possible. Companies that distribute AI applications may be required by governments to share their data; they are also run by capitalists, who themselves have political interests. Lastly, there may also simply be a data leak, as a result of hacking or an accident.

#### **3. Who is the target audience of the text, and what is the best way to reach them?**

Are they activists, or a public far removed from political organizations? A group of people who work in the same field, or different groups?

At this stage, it is useful not to be too economical with words when defining the target audience, even if this means having to narrow it down later on. Why? For example, if the only term used is “young”, AI can sometimes fall into the same traps as any communication that specifically targets a young audience: awkward references to video games or YouTubers.

We are all defined by more than just our age, so it is useful to combine several dimensions or target the most relevant ones in relation to the content of your text and where it will be published. Some – *mostly French* – examples: *“engineering students in Lyon”*, *“operators in an Amazon warehouse in*



*Montélimar*", "young urban parents", "Instagram users aged 18 to 30", "production workers in a Pasquier factory", "people in the Paris region who commute to work by RER", etc.

If the result goes too far in the opposite direction and becomes a little too specific, some of the details can be removed or reworded. As in cooking, it is a matter of finding the right balance of ingredients.

#### **4. What political viewpoint?**

As a general rule, your AI conversation partner will often be as consensual as possible. Asking a question without specifying a political angle will often result in receiving roughly the same answer you would get from a journalist working for a mainstream or public media outlet: the blame is shared, there are good solutions on both sides, etc.

If you are looking to generate text that should have a certain political or philosophical approach, you must indicate this in your prompt. This indication can be explicit (for example, by asking the AI to generate "*from a Marxist point of view*") or more indirect. Certain words or expressions will steer the generation toward a particular angle: "*comrade*", "*nationalization*", or "*self-organized collectives*" are all examples that carry their own political dimension and will more or less discreetly point in the direction you want to go.

Another problem that may arise when mentioning "*Communism*" or "*Marxism*" in the prompt is that some AIs may sometimes use language or examples taken directly from Marx's context (steam locomotives, emphasis on coal, etc.) or Lenin's. If this is not what you want, be creative in emphasizing our era: "*Marxism of 2025*", "*for a contemporary audience*", "*21st-century communism*", etc.

#### **5. Which format for the generated text?**

Based on the other elements (objectives, target audiences), you may already have a good idea of the format you want to generate.

Here are some ideas for the type of information that might be included: written or "*spoken*" text (for speeches, training sessions, etc.), intended for a specific platform (tweet, Facebook post, etc.), quantity to be generated ("*3 short paragraphs*", "*less than 300 words*"), mentions of specific elements ("*clear punchlines*", slogans), ...

Thanks to this preparation work, you have already started writing a large part of the content for the future prompt. Next, you need to decide on its structure!

### **Prompting techniques**

The techniques below are just a few of the many possible approaches we can take in our communications with AI, accompanied by examples of scenarios for militant use.

***No need to know how to code!***



None of these require programming knowledge, as they all use structures from our language that can sometimes be found elsewhere (role-playing, project management, logic and political thinking in general, etc.).

It is therefore only natural that some of them may seem familiar to you, and it is entirely possible that you may invent your own variations. This is even a desirable outcome.

### **Prompt RTF (Role Task Format) or CTF (Context Task Format)**

Behind these two somewhat complex acronyms lies a very straightforward technique that will encourage you to use the various elements you prepared before generation: the context, the task definition, and the output format.

#### **• RTF, or CTF?**

Same technique, two approaches: you can either define the role that the AI will take on to generate the text (that of an activist, a trainer, a researcher), or the context related to the task (organizing a training session, an event, etc.). You can even decide to combine the two, depending on what seems most natural for the task you want to tackle.

#### **• So, how does it actually work?**

Depending on the complexity of the prompt (particularly its length), we will be more or less formal in our use of this technique. For a prompt that can be stated in three simple sentences, we can just write one sentence for each aspect, such as:

*" You are a communist activist specializing in popularizing Marxism and contemporary social issues.[Role] Create a draft call to protest against student poverty in 2025.[Task] Your text should use an accessible but radical tone, and be structured in short paragraphs with subheadings.[Format]"*

*Answer of Mistral Medium 3.1 and GPT-4o mini in annexes 3 and 4.*

Here, each sentence corresponds to one of the dimensions of the prompt: the first assigns a role to "play" for the AI in its generation (communist activist), the second specifies the task to be accomplished (producing a call to protest), and the last adds formatting information (tone, structure).

The information in brackets serves as a guide for your reading; it is not necessary for such a short and clear prompt (each part is already delimited by a separate sentence).

**This first prompt is still very general and gives considerable creative freedom to the AI.** It will draw on its own elements to demonstrate student poverty and may even propose its own demands!

Neither will be the result of original political thinking on the part of the AI tool, but rather what it calculates to be the closest imitation of this type of text. It is therefore a possible starting point in terms of shape (if it includes

elements that are relevant to your objectives), but never a final text that will be generated in this way.

### Chain Of Thought (COT)

The Chain Of Thought method, as its name suggests, guides AI towards a series of interconnected ideas. In other words, it breaks down text generation into several steps that we want to see in the final text. To do so, we will simply add them directly to the prompt.

**This method is ideal for complex tasks**, such as analytical work, which can take many different paths. It allows you to maintain a certain degree of control over which paths will ultimately be taken.

Here's an example:

*Suggest a campaign strategy to achieve free public transportation in a metropolitan area.*

1. **Analyse** first the economic arguments (social cost of private cars, existing public subsidies) and ecological arguments (reduction in emissions).
2. **Identify** the stakeholders to be mobilized (users, unions, associations) and the political levers (referendums, pressure on elected officials).
3. **Propose** a calendar of actions (demonstrations, petitions, symbolic occupations) inspired by similar campaigns in Europe.

Here, parentheses are used to include quite a few details about what each point may correspond to (ecology in point 1 focuses on emissions, for example). Depending on the tasks, you need to find a balance between what you want to leave open for the generation and what needs to be partly predetermined and more predictable.

It is even possible to combine this method with the previous one. Returning to our example of student poverty and making its various elements more visible, we could end up with a prompt like this:

**Role** : You are a communist activist specializing in popularizing Marxism and contemporary social issues.

**Task** : Create a draft call to protest against student poverty in 2025. Your text must include:

1. A simple explanation of the structural causes of student poverty (**include causes here**);
2. A brief recap of recent actions on this issue (**include relevant actions here**);
3. Concrete policy proposals (**include your proposals here**).

**Format** : Your text should use an accessible but radical tone, and be structured in short paragraphs with subheadings.

## Few-Shot prompting

Basically, the idea behind this technique is to “feed” the AI with several examples of the type of results you want to achieve.

If you are already accustomed to writing texts, this may include relevant excerpts from your own work, with instructions to reproduce the same writing style, line of argument, or format on a new topic.

Some examples of possible uses: teaching an AI – *which has access to the internet* – to reproduce the style of a particular press review, adapting a new text to the editorial style of a platform, etc.

## When one prompt is not enough

Sometimes, the task is too complex to be dealt with adequately in a single text generation. Several methods can then be considered.

### Decomposed Prompting

If the problem being addressed has many dimensions, it may be useful to simply break it down. This is, in a nutshell, the approach of *decomposed prompting*.

Take, for example, organizing a series of Marxist lectures on a campus. How should the organizational issues that arise be divided up? The following dimensions or areas could be considered:

- **Content:** definition of the specific theme, identification of speakers and possible topics for presentations.
- **Logistics:** venue management, necessary equipment, travel arrangements and accommodation for non-local speakers, financial aspects.
- **Communication:** what campaign on campus, social media? Potential partners. Promotion of the content created after the conferences.

You can list in each of these areas the questions that are currently unanswered and which might be addressed in a prompt. The end result is a kind of organizational plan, with each sub-section including prompts where necessary.

After generating and saving the corresponding answers for each section, the result you have is a collection of specialized “expertise” that goes beyond the information that you could get from a single general prompt.

### Self-Reflection prompt

Very straightforward method:

1. Order a first text-generation for any task you need to address, using the method that suits you best.
2. In the same conversation, ask the AI to generate a critique of its own text, either general or based on a specific criteria of your choice. For example, “*Write a critique of your text, focusing on its accessibility to an audience unfamiliar with activism.*”



3. Ask the AI to build on this critique to generate a new version of the first text.

This method relies on AI's ability to process information from texts it has already produced (in an ongoing conversation) to generate subsequent ones; in this way, it can mimic critical thinking.

## What Next?

- **Critical assessment**

After generating the text, it's time to revisit your preparatory questions and use them as critical tools. Does the text enable you to achieve your defined objectives? Does it seem to be the right way to address your target audience, or does it contain any awkward phrasing? Are the ideas being put forward the right ones?

Re-examine your prompt in light of the result: did any of your words have too strong an influence?

- **When is the right time to stop generating?**

If the result is far from expectations, it is possible to revise the composition of the prompt with other words, or try a new technique to achieve a better result.

On the other hand, if it only contains a few minor errors or issues after several attempts, you can consider it a valid starting point that you can now modify and improve. The advantage of using AI tools is that they allow you to focus on the structure and ideas of a document, but it is unlikely that they will completely eliminate the work of writing itself.

- **Save your prompts**

If your text generation meets your expectations, it is useful to keep a copy of the prompt, along with information about the type of AI used (and, if possible, its version), and perhaps even the generated text itself. If the task that this prompt addresses is common to other activists, why not share it?

- **Facts, figures and statistics**

Never trust any information generated by AI without verifying it. Even plausible elements may be slightly or entirely false; it is in the nature of this tool to offer information that seems probable, with a certain degree of certainty.

Using artificial intelligence in your areas of expertise can help you move very quickly, as inconsistencies are easy to spot in these situations; in other areas, be cautious. Here are a few techniques to adopt:

### 1. Request sources

If an AI has access to Internet, it is possible to ask it to link its statements to sources. Do not hesitate to do so in the prompt when relevant. Be specific in your requests. What type of sources meet your needs? (Scientific research portals, certain types of online media, specific authors, etc.)

## 2. Test the links

When a link is provided as a source, it may not lead anywhere. This could be an indication that it has been "*made up*" and that the associated figure or fact may be false. Do not take the presence of a link as sufficient guarantee: visit them.

## 3. Quick-checking links tip

When the link works, reading an entire page to verify information partly defeats the purpose of saving time. But all is not lost! If you are looking for a number, date, or a name (which fits most cases), you can quickly search the web page or PDF document to find the relevant excerpt(s) (shortcut keys Ctrl and F on most browsers).

## 4. Ask follow-up questions

Sometimes it is not possible to obtain proof in the form of a link. For example, because the AI has had access to copyrighted content and part of its code prevents it from communicating about it, or simply because you are using a service that cannot access Internet.

In such cases, you can ask questions that will help you get a clearer picture, for example: "*Can you give me some concrete examples or practical cases that illustrate this statement?*", "*Are there any contradictions or debates surrounding this information?*", "*Suggest a way for me to verify your statement.*"

## 5. Cross-check sources

If you still have doubts, you can also verify certain information through authoritative sources in the relevant field. The keywords used by AI in its response may sometimes be the same as those you would use to conduct your own research.

- **Calculations: Opt for a good old calculator**

Unfortunately, the effectiveness of calculations depends on the context, prompt, the AI model, and it is probably safer not to trust the result of a calculation that you cannot verify. ***This applies in particular to all statistics calculated in a generation from external sources, even if these sources are reliable..***

Be especially careful with tables that summarize and mix figures from different units and sources in order to draw conclusions. For conversions from one unit to another, many online services are more efficient, and for everything else, a calculator remains a more reliable tool.

**Ironically, if you're not great at math, AI can be a good teacher** and explain the stuff you're struggling with in a way that's as easy to understand as you need it to be. For example, you can use it to learn a simple method for calculating percentages or proportions, do cross-multiplication, or even more advanced uses like calculating statistical correlations, explaining algebra concepts, and so on. This is a good attitude to adopt in general: don't be dependent on AI's answers, but use it to learn the knowledge you lack so that you can judge its answers, even those you think are beyond your reach.

### • **Language**

The more widely a language is used internationally (especially online), the more likely it is that large AI models will represent it accurately in their text generation. However, these models are still primarily designed using English texts, so when generating non-English texts it is important to be aware of the occasional invented word or over-representation of unusual phrasing. That said, errors of this kind are becoming increasingly rare.

#### • **AI can sometimes be a bit too friendly**

If you ask for advice from AI in a conversational tone, chances are high that the result will be at least partially positive.

Sometimes, simply mentioning that the text you are using in the prompt is your own will generate praise. Your questions and comments will always be relevant or excellent, even when they may contain errors of understanding in their wording. This is normal: if you speak to the AI in a polite tone rather than as if it were a machine, it will be statistically encouraged to generate a response in the same spirit. This is nice, but sometimes we need a little more critical feedback!

When you do want to receive a more nuanced opinion, you need to positively influence the generation in that direction, for example with sentences such as:

*"If you find any errors (logical, comprehension) or mistakes in my text, mention them and justify your selection."*

#### • **It's you who evaluate AI, not the other way around**

Maybe AI makes fewer spelling mistakes or uses more elegant phrasing than you do, but it has no real understanding of the text it produces, of politics, or even of our world.

You are therefore much more qualified to judge its work than the other way around. It may be useful to ask the AI for corrections or modified versions of a text, but decisions regarding the organization of your work and final approvals should always remain yours.



## Wrap-up

*by Hugo Pompougnac*

You have reached the end of this short guide (excluding the annexes). We do hope you enjoyed reading it, or at least found it useful. If so, we have a favor to ask of you. As mentioned in the introduction, this version of the guide is still very much a work in progress. We believe there is plenty to add, perhaps remove, correct, and improve. With this in mind, our primary concern is its usefulness to activists, so your feedback after reading it is very important. Please email us at [contact@espaces-marx.eu](mailto:contact@espaces-marx.eu) with any comments, criticism, or suggestions.

You can also write to us if you would like to actively contribute to the writing process. If you are comfortable with digital tools, we invite you to interact with us via GitHub, visit the guide repository ([github.com/espaces-marx/ai-radicals](https://github.com/espaces-marx/ai-radicals)), fork it, and send us a Pull Request. We welcome all contributions and would be happy to build a larger activist community, working together to develop skills so that the left can better master new technologies.

For the moment, our intention for the next version of the guide is to:

- Provide concrete examples and in-depth explanations of “advanced” AI tools (agents, connectors, libraries, etc.).
- Detail other tools; for now, we have focused on Mistral, although much of our advice applies equally well to other AIs.
- Provide more specific advice on AI-assisted data synthesis (summarizing a legal text, etc.).
- Teach AI-assisted graphic design techniques: not just by writing a trivial prompt and marveling at the result, but by combining AI with other techniques to create a poster, sticker, Instagram visual, etc.

Looking ahead, we would like to explain (in the clearest way possible) how to retrain an AI model at home to transform its purpose. But that's a more geeky goal; for the time being, we believe there are more pressing matters.

If you would like to comment on these objectives or suggest others, please contact us at our email address.

*Arrivederci!*



## Annexes

*All annexes were initially generated in French, then translated in English for the purpose of this guide. You might not get the same result, even with the same prompt and AI model.*

### Example Leaflet

*(Based on the press release)*

---

#### MACRON OUT, POWER TO THE PEOPLE!

**Why join the protest on September 10?** The government wants us to pay for the crisis: delayed retirements, frozen wages, dismantled public services. But anger is rising, and this day could change everything.

##### 3 reasons to take to the streets:

- Shift the balance of power: The Yellow Vests proved it—when we stand united, those in power back down.
- Our lives matter more than their profits: Billionaires are feasting while our pensions vanish. Enough is enough!
- Another society is possible: September 10 is our chance to show that social revolution is on the agenda.

Meet up: **2 PM, Place de la République** (Paris)

Protests across France (check locally for details)

***"They have the billions, we are the millions. On September 10, let's shake the system!"***

---

*[Customize with your logo, emojis and contact info]*

### Example Script

*(Based on the editorial "September 10: Revolution on the Agenda")*

---

#### [Introduction – 0:00 to 0:20]

"September 10 isn't just another protest. It's a test: Do we accept defeat, or do we fight to win? I choose to fight. And you?"

#### [Development – 0:20 to 1:30]

"For years, they've been tightening the screws: retirement at 64, hospitals on their knees, skyrocketing rents. Meanwhile, shareholders are breaking profit records. We can't take it anymore.



On September 10, we show them we won't back down. Just like in 2023, when strikes forced the government to retreat on pensions. Proof that when we stand together, we can win.

But this time, it's not just about stopping a reform. It's about overturning the whole system. Because this system is rotten to the core."

#### **[Conclusion – 1:30 to 2:00]**

"So be there on September 10. In the streets, with your friends, coworkers, and neighbors. Because revolution isn't built alone—it's built block by block, protest after protest.

In Paris, Lyon, Marseille: we'll be thousands strong. And this time, we won't go home without victory.

On September 10, revolution is on the agenda. And it starts with you."

---

*[Adapt with protest footage, subtitles, and a powerful soundtrack.]*

### **Example of RTF generation with Mistral Medium 3.1**

**Call to action: Let's fight together against student poverty!**

---

#### **Poverty is not inevitable; it is a political choice!**

In 2025, studying is still synonymous with hardship. Far from being a mere transition, student poverty has become the norm: exorbitant rents, exhausting part-time jobs, insufficient grants, and crowded lecture halls where we are prepared for an increasingly precarious job market. Meanwhile, the profits of large corporations and tax breaks for the wealthiest are skyrocketing.

**Our hardship finances their prosperity.**

#### **Broken promises, violated rights**

We are promised social mobility, but what we get are crumbling staircases. Successive governments have abandoned students: disguised increases in tuition fees, cuts to financial aid, and a public university system that is being suffocated. Meanwhile, unpaid internships are on the rise, and fixed-term contracts follow one after another with no hope of stability. **Our future is being sacrificed on the altar of profit.**

#### **The anger is legitimate, the fight is necessary**

Given this situation, resignation is not an option. All over France, collectives are rising up: university occupations, rent strikes, university cafeteria blockades. These actions show one thing: **when we organize, we can bring the system to its**



**knees.** Precariousness is not inevitable, but the result of a capitalist system that exploits us. It's time to fight back.

### Join the movement!

**Saturday, October 4, 2025, let's take to the streets en masse!**  
Meet at 2 p.m. at Place de la République (Paris) and in all major cities across France.

### Our revindications :

- **Free tuition:** not a single euro more for enrollment fees!
- **A decent student income:** a minimum of €1,200 for everyone!
- **Social and affordable housing:** stop real estate speculation at our expense!
- **Stable and decent jobs:** an end to precarious internships and fixed-term contracts!

**Let's organize, unionize, and fight together!** Precariousness is a weapon used against us. **Let's turn it into collective strength.**

**Capitalism is stealing our future. Let's take it back by force.**  
#StudentPoverty #GeneralStrike #FreeEducation

---

## Example of RTF generation with GPT-4o-mini

---

### Together against student poverty

Student poverty is a scourge that affects millions of young people in France. Every day, students struggle to survive, juggling classes, part-time jobs, and exorbitant rents. It's time to put an end to this unacceptable situation!

### A system on its last legs

The current economic model prioritizes profit at the expense of education. Tuition fees are skyrocketing, financial aid is insufficient, and living conditions are deteriorating. Students should not have to choose between eating and succeeding in their studies. We must demand access to education for all, without conditions.

### Solidarity, our weapon

In the face of this precariousness, solidarity is our best ally. We must unite – *students, workers, and citizens* – to make our voices heard. Past struggles have shown us that collective mobilization can bring about change. Together, we are stronger!

### Call to action



We call on you to join the demonstration on [date of demonstration] at [location of demonstration]. Together, let's make our anger and demands heard. Let's demand concrete measures: increased scholarships, rent caps, and free access to healthcare.

**For a better future**

Let's not allow uncertainty to define our future. Let's fight for accessible, high-quality education, for a world where every student can thrive without fear. Change starts now, and it starts with us. Let's come together in large numbers and be determined!

---