

How to use DeepSeek AI

A simple guide to learn how to use the LLM models from DeepSeek

Artificial Intelligence



by Pallav Pathak
January 23, 2025

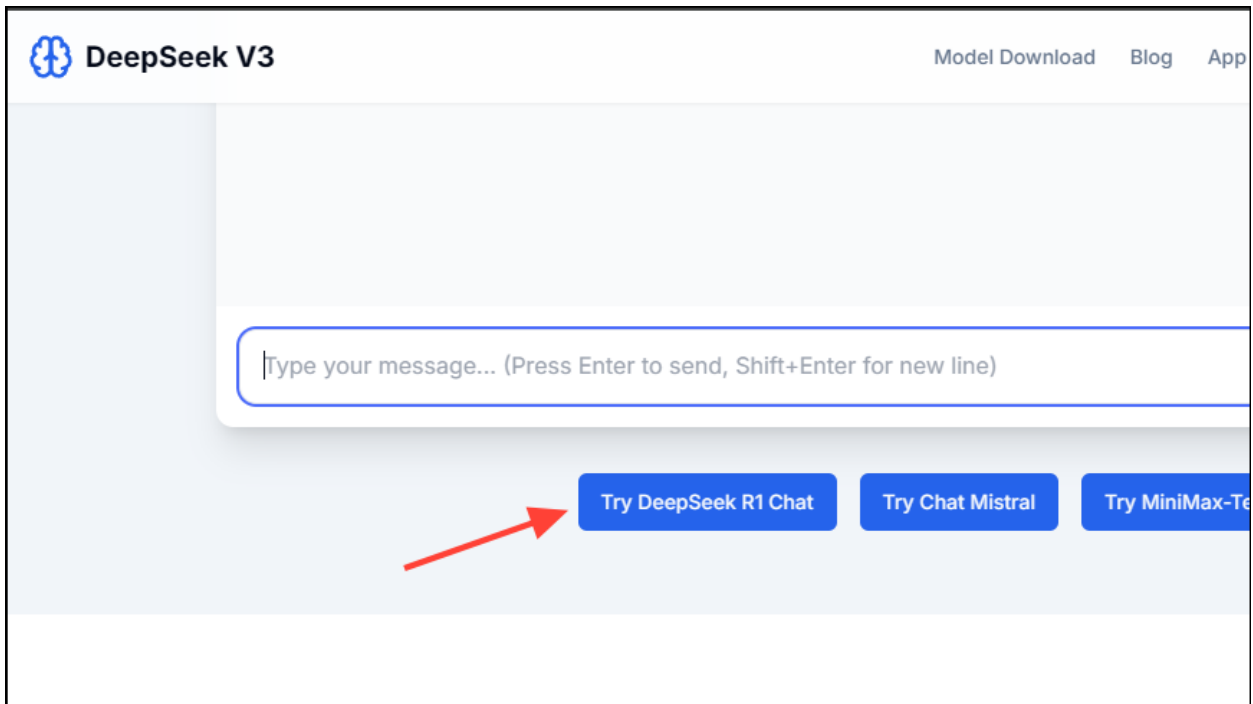
DeepSeek is perhaps best known as the Chinese startup responsible for developing the DeepSeek V3 AI model. This is an open-source, scalable model that can process data at incredibly high speeds and has very advanced reasoning and understanding capabilities. More importantly, it has outperformed other more famous models like GPT-4o, Qwen 2.5 Coder, and Claude 3.5 in tests.

Now, DeepSeek has launched two new AI models, DeepSeek R1 and DeepSeek R1 Zero, which can match the performance of OpenAI's o1 model and are much more affordable.

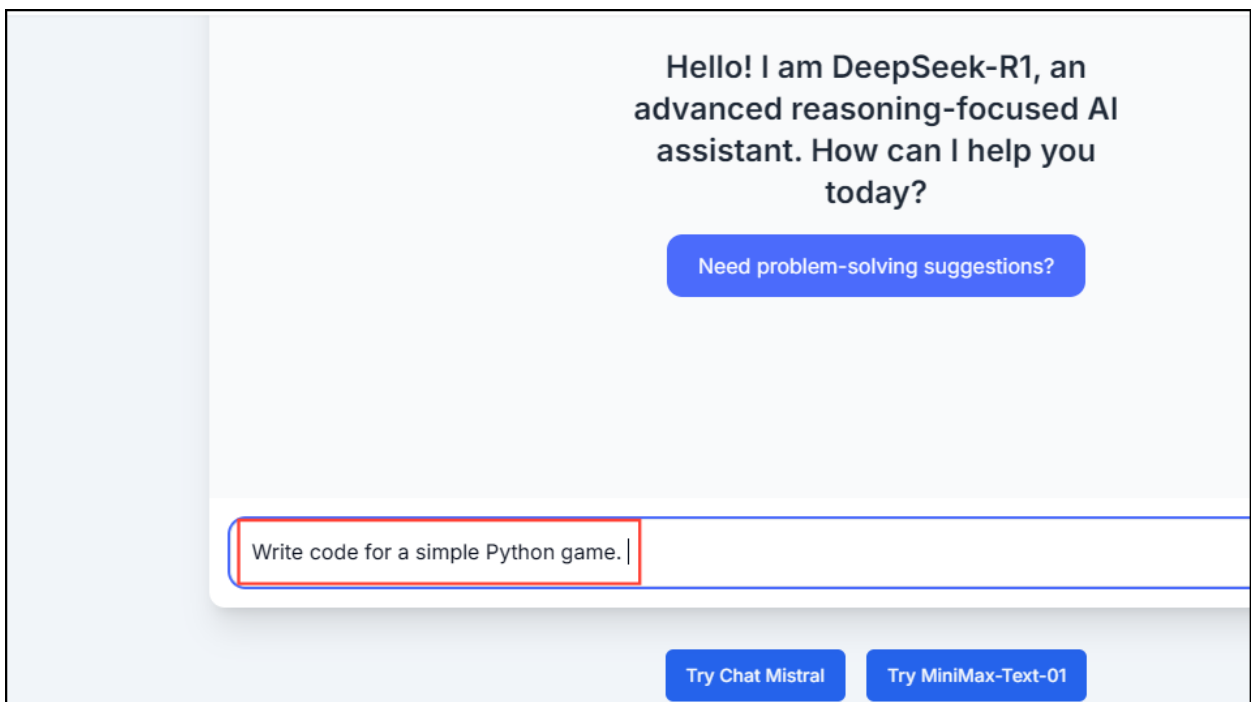
You can try out DeepSeek AI on your computer without having to purchase a subscription plan, though a subscription is required if you want to use the advanced features of various DeepSeek models. DeepSeek models are also available on HuggingFace.

Accessing DeepSeek R1 through a web browser

1. Open a web browser on your computer and head to www.deepseekv3.com
2. Once the website opens, click the 'Try DeepSeek R1 Chat' button below the chat box.



3. Now, you will see the DeepSeek R1 Chat interface on your screen. You can ask the AI questions by typing them in the chat box.



4. Once you've typed in your query, hit Enter or click on the 'Send' button and the AI will provide you with a response.

Here's a simple "Guess the Number" game in Python:

```
import random

def guess_the_number():
    # Generate random number between 1 and 100
    secret_number = random.randint(1, 100)
    attempts = 0

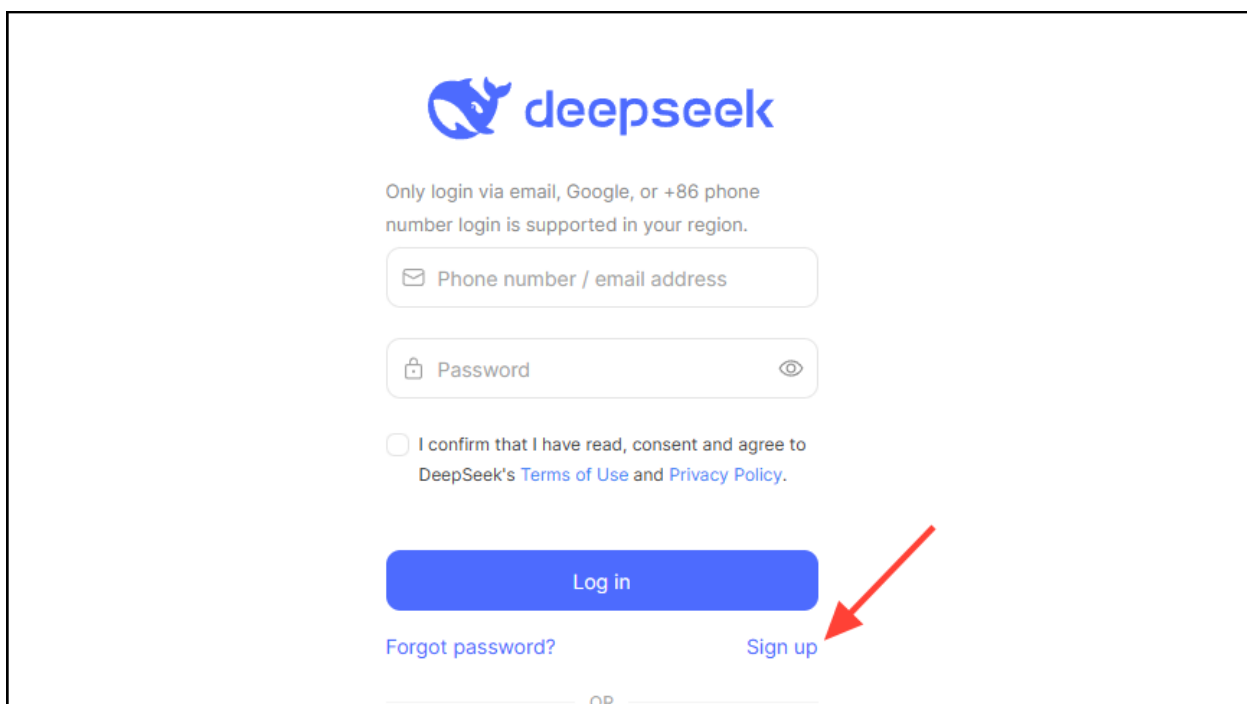
    print("Welcome to Guess the Number!")
    print("I'm thinking of a number between 1 and 100.")

    while True:
        try:
            guess = int(input("Your guess: "))
            attempts += 1

            if guess < secret_number:
                print("Too low! Try again.")
            elif guess > secret_number:
                print("Too high! Try again.")
            else:
                print("Congratulations! You guessed the number in", attempts, "attempts.")
                break
```

Accessing DeepSeek V3 Coder via API

1. Open your browser and go to chat.deepseek.com.
2. Click on the 'Sign Up' option on the page and create your account.



deepseek

Only login via email, Google, or +86 phone number login is supported in your region.

Phone number / email address

Password

I confirm that I have read, consent and agree to DeepSeek's [Terms of Use](#) and [Privacy Policy](#).

Log in

[Forgot password?](#) [Sign up](#)

OR

3. Once you create an account, you should get an API key.

4. Go to python.org and download Python from the website if you haven't already installed it. Make sure to check the box for 'Add python.exe to PATH' during installation, otherwise, you will have to navigate to the Python directory whenever you need to run a Python command.
5. After installing Python, you can install the required SDK. Since DeepSeek API relies on OpenAI-compatible API formats, you can use it with the OpenAI SDK or software compatible with OpenAI. Open the Command Prompt and type `pip install openai` and hit Enter.
6. After setting up your development environment, set the base URL to `https://api.deepseek.com` and configure your API access.
7. Now you can access DeepSeek AI by making API calls.

Deploying DeepSeek V3 locally

If you want to deploy DeepSeek AI locally, you will need to set up the required environment for it and download the local models. Keep in mind that local deployment is best suited for Linux distros like Ubuntu, not for other operating systems like Windows. So, you will need to create an environment similar to Linux in Windows if you want to deploy DeepSeek locally. To deploy DeepSeek locally, you will need a GPU with CUDA support, Python version 3.8 or higher, at least 16 GB of RAM, and CUDA and cuDNN.

1. To deploy DeepSeek locally, install Windows Subsystem for Linux on your system.
2. Clone the DeepSeek repository using the following command: `git clone https://github.com/deepseek-ai/DeepSeek-V3.git`.
3. Go to the inference directory and install the required dependencies. You can find them in 'requirements.txt'. Enter the following commands to install the dependencies: `cd DeepSeek-V3/inference` and `pip install -r requirements.txt`.
4. Use the following command to convert the model weights to a specific format:

```
python convert.py --hf-ckpt-path /path/to/DeepSeek-V3 --save-path  
/path/to/DeepSeek-V3-Demo --n-experts 256 --model-parallel 16 .
```

5. Now you can chat with DeepSeek.

Things to know

- You can access certain DeepSeek models for free but will need to pay for others, with the costs depending on your usage. For instance, DeepSeek-V3 has a pricing structure of \$0.14 per million input tokens and \$0.28 per million output tokens. And, as mentioned before, if you want to access advanced options or premium features, like customizing your preferences or saving your chat history, you will need to create a DeepSeek account on the website.
- Since the DeepSeek R1 model can solve complex mathematical and reasoning problems, you can use it for tutoring and education purposes, debugging, code generation, etc.
- DeepSeek has also released DeepSeek R1 distilled models, which are more efficient, smaller versions of the model, and can run on consumer devices.
- The R1 model also combines real-time web browsing, allowing you to search for information online and analyze it while combining it with the data the model already has, which makes DeepSeek AI incredibly powerful and similar to other models like Perplexity AI.

[NEWER POST](#)[OLDER POST](#)



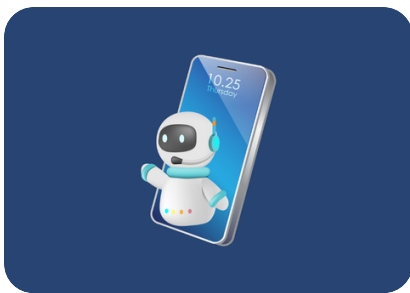
How to speed up slow charging on a laptop

DeepSeek releases open-source reasoning model, DeepSeek-R1, to rival OpenAI's o1



Member discussion

YOU MIGHT ALSO LIKE



How to use Perplexity Assistant on Android



DeepSeek releases open-source reasoning model, DeepSeek-R1, to rival OpenAI's o1



What is Scribe AI and how to use it

NEWSLETTER

Get all the latest posts delivered straight to your inbox.


Your email address

Subscribe



Technology explained, simply and clearly.

Social

 Facebook

 Twitter

Links

Contact

About

Privacy Policy

Terms & Conditions

Links

Sign up

© All Things How 2025.