

Build an AI bedtime storyteller with Stable Diffusion



Sean McManus

Author of *Mission Python*, *Web Design in Easy Steps*, and *Raspberry Pi For Dummies* (with Mike Cook). Get free chapters at Sean's website.

sean.co.uk

You'll Need

- ▶ ChatGPT account and API key platform.openai.com
- ▶ Pimoroni HyperPixel Square (optional) magpi.cc/hyperpixel4
- ▶ External speakers (optional)

Tell Bedtime Stories who's the hero and where it happens, and the program creates an illustrated tale for you, using the magic of artificial intelligence

Once upon a time there were two artificial intelligences, who were famed throughout the land. ChatGPT was an inventive storyteller, while Stable Diffusion was an ace illustrator. Together, they made up children's stories, which Raspberry Pi read aloud while showing the pictures. Listeners could put themselves, or their friends, in the stories, and decide where they should take place. Sometimes the stories were strange, and the pictures even more so. But as an art experiment, Bedtime Stories was a revelation. Anyone spying on the code could discover how to use ChatGPT and Stable Diffusion to create text and pictures from Python.

01 Get your DeepAI API key

DeepAI aims to make artificial intelligence (AI) accessible to new developers. The company provides an online text-to-image tool, accessed through an application programming interface (API). You send it a prompt (a request) and it sends you back an image. The API has been around since 2016 and its latest version uses the Stable Diffusion image generation software. Visit deepai.org to create your account. You can buy

credit for 100 API calls for \$5. On your Profile page, change the settings to allow unfiltered content. The filtered API rejected requests that included terms such as 'hand in hand' or 'creepy', resulting in incomplete stories.

02 Install Python modules

We'll need three Python modules for this project. OpenAI's ChatGPT is used to create the story text and the image prompts for DeepAI. An image prompt is a description of the image we'd like the AI software to make for us. The `requests` module is used to access Stable Diffusion through the DeepAI API. The `pyttsx3` module reads the stories aloud. Using speech in this project makes it feel like a friendly robot is reading you a bedtime story, but also means we can dedicate the screen to the story images. Open a terminal window and enter these instructions to install the Python modules and text-to-speech software. You may see some warnings about OpenAI not being on PATH, but you can safely ignore them.

```
pip install pyttsx3 requests openai
sudo apt install espeak mopidy
```

03 Create the ChatGPT prompt


This project has two parts: Bedtime Story Maker (**story_maker.py**), and Bedtime Story Reader (**story_reader.py**). Stories are created and saved by the Maker and loaded and read by the Reader. Take a look at **story_maker.py**. When you run the program, it asks you to enter the character names and types, location, and genre for the story. It then builds them into a string called **story_prompt** that is used to ask ChatGPT to make the story. The story prompt also asks for an image prompt for every paragraph, starting with 'Image Prompt:'. The story prompt tells ChatGPT to suggest a book illustrator for each picture.

04 Using the ChatGPT API

Lines 33 to 39 send ChatGPT the story prompt. Add your own ChatGPT API key. We've told ChatGPT it is a children's author. Hints like this help to guide it to the best responses, based

“ Saving the stories means you can play them on demand, without waiting for them to generate ”

on the vast amount of online content it's been exposed to. The response comes in as a single piece of text, which should alternate between story



▲ ChatGPT came up with the idea for this female ghost with long, flowing hair and ethereal robes. Stable Diffusion made it real



paragraphs and image prompts. Lines 42 to 46 split that response into a list of image prompts and story paragraphs. Sometimes ChatGPT fails to return any image prompts, so the **while** loop keeps asking ChatGPT for stories until it also sends some image prompts.

05 Saving the story and image prompts

The story paragraphs and image prompts are saved to text files using the **write_list()** function. It accepts a list and a file name and saves each list item using the file name with a number on the end (e.g. **story-0.txt**, **story-1.txt**). They're saved in a folder named with the time, from the year to the second, such as 'Story 2023-04-10-09-34-52'. Saving the stories means you can play them on demand, without waiting for them to generate, as often as you like. It also enables you to study what worked well, and what didn't. ChatGPT wrote the first version of this function for us, by the way. It's good at basic Python programming, as long as you know how to guide it. We used the online front-end to ChatGPT for that at **chat.openai.com**.

06 Generate Stable Diffusion images

Next, the program goes through the list of image prompts and sends each one to DeepAI's API. You'll need to add your own DeepAI API key in line 59. To make the images more consistent, the program adds something to each prompt to clarify who the characters are and where the story

Top Tip


Watch your wallet

The DeepAI API requests are inexpensive, but if you make a lot of stories, it adds up fast.



Alert!
Strange images

Stable Diffusion may produce images that are unsettling for young children, such as distorted faces.



▲ The AI image generation software does particularly well with sci-fi images

takes place. We've also added a negative prompt to provide some hints that can help improve image quality. We asked for images that are 720x720, to match the size of the HyperPixel display you are using. Valid values are between 128 and 1536, but you might get odd results below 256 and above about 700. DeepAI returns your image as a URL, so the program downloads the file and saves it for you. It also displays the URL, so you can download the image (you've paid for!) manually if something goes wrong. You can find a simple version of the code to generate and download an image at magpi.cc/seandeepai.

Top Tip

Flip the images

If your display is upside down, you can flip your images by inserting this before line 29 in Story Reader:

```
image_to_show
= pygame.
transform.
flip(image_
to_show, True,
True)
```

07 Do your quality control

Use the Run button in Thonny to run **story_maker.py** a few times to make some stories. When it finishes, the story folder should contain story files (e.g. **story-0.txt**, **story-1.txt**); image prompt files (**prompt-0.txt**, **prompt-1.txt**); and an image file for each prompt (**1.jpg**, **2.jpg**, etc.). The Reader reads you the story paragraph aloud while it shows you the related image, so there should be the same number of image and story files. Sometimes there's an extra story file because ChatGPT adds a title, and the Reader can handle that. Occasionally, ChatGPT doesn't send enough image prompts, so there aren't enough pictures. In that case, you can still read the story

and enjoy the illustrations yourself, but the Reader won't work. Find a valid story folder you want to play and copy its folder name.

08 Set up Bedtime Story Reader

Now you've made some stories, let's use the Bedtime Story Reader (**story_reader.py**) to enjoy them. Add the name of the folder where your story is stored in line 3 of the program.

09 Discovering the story files

After it's set up the Pygame window for displaying images, and the voice for reading them, the program indexes the story and image files. Their file names go into the **story_files** and **image_files** lists, which are then sorted into the correct order. If there are more story files than image files, the **story_has_title** variable is set to True.

10 Displaying the picture

The program goes through the **story_files** list and displays the associated image for each story paragraph. If the story has a title, the folder will be one image file short because there is no image prompt for the title. In that case, the first image needs to be shown for both the title and the first paragraph. Line 24 ensures this happens.



▲ It's common for characters in AI-generated images to have odd-shaped heads, distorted faces, and an unconventional number of limbs. It all adds to the surreal effect

story_maker.py

> Language: **Python**

**DOWNLOAD
THE FULL CODE:**



magpi.cc/bedtimestories

```

001. # Story Maker - by Sean McManus - www.sean.co.uk
002. import os, openai, requests, datetime, random,
    time
003.
004. def write_list(list_to_write, filename):
005.     for number, paragraph in enumerate(
        list_to_write):
006.         filepath = os.path.join(directory,
            f"{filename}-{number}.txt")
007.         with open(filepath, "w") as file:
008.             file.write(paragraph)
009.
010. date_time_str = datetime.datetime.now().
    strftime("%Y-%m-%d-%H-%M-%S")
011. directory = f"Story {date_time_str}"
012. os.mkdir(directory)
013. image_prompts, story_paragraphs = [], []
014. character1_name = input(
    "What is your main character's name? ")
015. character1_type = input(
    "What kind of a character is that? ")
016. character2_name = input(
    "What is your second character's name? ")
017. character2_type = input(
    "And what kind of a character is that? ")
018. venue = input(
    "Where does the story take place? (e.g. in a
    castle, on Mars) ")
019. genre = input("What is your story genre? ")
020. story_prompt = f>Please write me a short {genre}\
021. story. In this story, {character1_name} is a \
022. {character1_type} and {character2_name} is a \
023. {character2_type}. The story takes place
    {venue}.\
024. For each paragraph, write me an image prompt for\
025. an AI image generator. Each image prompt must\
026. start in a new paragraph and have the words\
027. 'Image Prompt:' at the start. Choose a book\
028. illustrator and put something in the image\
029. prompts to say the images should be made in the\
030. style of that artist."
031. while len(image_prompts) == 0:
032.     print("Making ChatGPT request")
033.     openai.api_key = "YOUR API KEY" ### PUT YOUR
        API KEY HERE
034.     ChatGPT_output = openai.ChatCompletion.
035.         create(
036.             model="gpt-3.5-turbo",
037.             messages=[
038.                 {"role": "system",
039.                 "content": "You are a children's author."},
040.                 {"role": "user",
041.                 "content": story_prompt}
042.             ] )
043.     new_story = ChatGPT_output.choices[0].
        message["content"]
044.     print(new_story)
045.     for paragraph in new_story.split("\n\n"):
046.         if paragraph.startswith(
047.             "Image Prompt"):
048.             image_prompts.append(paragraph)
049.         else:
050.             story_paragraphs.append(paragraph)
051.     write_list(story_paragraphs, "story")
052.     write_list(image_prompts, "prompt")
053.
054. for number, image_prompt in enumerate(
    image_prompts):
055.     image_prompt += f"{character1_name} is
    {character1_type} and {character2_name} is
    {character2_type}. They are {venue}."
056.     print(f"Generating image {number}")
057.     r = requests.post(
058.         "https://api.deepai.org/api/text2img",
059.         data={'text': image_prompt,
060.             'negative_prompt': "poorly drawn
061.             face, mutation, deformed, distorted face,
062.             extra limbs, bad anatomy",
063.             'width': "720", 'height': "720",
064.             'grid_size': "1"},
065.         headers={'api-key': 'YOUR API KEY'} ###
066.         PUT YOUR API KEY HERE
067.     )
068.     image_url = (r.json()["output_url"])
069.     print(f"Image {number} is at {image_url}.
    Saving now...\n\n")
070.     filename = f"{number}.jpg"
071.     filepath = os.path.join(directory, filename)
072.     img_data = requests.get(image_url).content
073.     with open(filepath, 'wb') as handler:
074.         handler.write(img_data)
075.     print(f>Your files are in {directory}.")

```

story_reader.py

DOWNLOAD
THE FULL CODE:

► Language: Python



magpi.cc/bedtimestories

```

001. # Story Reader - by Sean McManus - www.sean.co.uk
002. # Put your story folder below
003. path = "Story 2023-04-10-12-35-04"
004. import os, pygame, pyttsx3
005. win_width, win_height = 720, 720
006. pygame.init()
007. windowSurface = pygame.display.set_mode((
    win_width, win_height))
008. pygame.mouse.set_visible(False)
009. voice = pyttsx3.init()
010. voice.setProperty('rate', 170)
011.
012. story_files, image_files = [], []
013. for file in os.listdir(path):
014.     if file.lower().endswith('.jpg'):
015.         image_files.append(file)
016.     elif file.lower().startswith('story'):
017.         story_files.append(file)
018. story_has_title = len(story_files) > len (
    image_files)
019. story_files = sorted(story_files)
020. image_files = sorted(image_files)
021.
022. for number, story in enumerate(story_files):
023.     if story_has_title:
024.         image_path = os.path.join(
    path, image_files[max(0, number - 1)])
025.     else:
026.         image_path = os.path.join(
    path, image_files[number])
027.     image_to_show = pygame.image.load(image_path)
028.     image_to_show = pygame.transform.scale(
    image_to_show, (win_width, win_height))
029.     windowSurface.blit(image_to_show, (0,0))
030.     pygame.display.update()
031.     story_path = os.path.join(path, story)
032.     with open(story_path, "r") as file:
033.         story = file.readlines()
034.     voice.say(story)
035.     voice.runAndWait()
036. pygame.quit()

```



▲ Here, Stable Diffusion has created a bright, colourful image that matches the look of a preschool book

It turns the list indices 0,1,2,3... into 0,0,1,2... and ensures we don't run out of images. If we had four story paragraphs (numbered 0 to 3) and three images (0 to 2), the last image and the last paragraph would coincide as they should. Images are scaled to fit the display. Although images are requested at 720×720, they come through at 768×768 because image sizes are accurate to the nearest multiple of 64.

11 Read out the story paragraph

Lines 32 to 35 open the text file for the story paragraph, read it in to the variable `story` and then use the text-to-speech module to read it aloud. For tips on customising the voice, see Raspberry Radio in issue 122 (magpi.cc/122).

12 Run the code

You can run Bedtime Story Reader using any display and in the desktop. Here we are using the HyperPixel Square display. To make the images fill the display, we boot to the command line and run the program using:

```
python story_reader.py
```

There is lots you can do to build on this project. You could add a menu for selecting stories, transitions between images or background music. The code shows you how to use AI to make text and images for any project, opening up a world of adventures. 🚀