

从语音转图像说起，生成式任务的前世今生

付杰

Jina AI 高级算法工程师



CONTENT

01 生成式任务简介

02 语音转图像任务

03 使用 Jina 快速搭建工作流

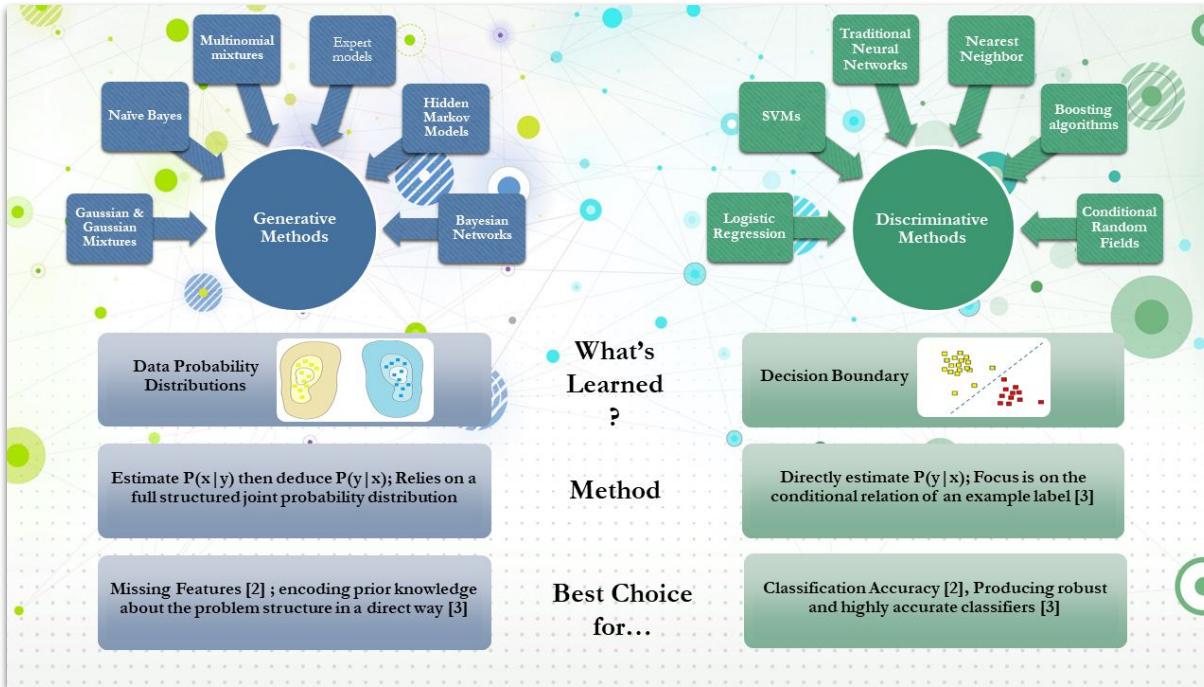
04 Demo 上手

生成式任务简介

01

生成式任务简介

- 判别式任务：分类 / 回归，从现有数据出发进行分析
- 生成式任务：学习已有的数据分布，生成全新的数据



生成式任务简介



Can you write a function that takes a string input and capitalizes the first letter of each word?

Sure, here is a possible function that capitalizes the first letter of each word in a string input:

```
def capitalize_first_letters(string):
    words = string.split()
    capitalized_words = []
    for word in words:
        capitalized_word = word[0].upper() + word[1:]
        capitalized_words.append(capitalized_word)
    return " ".join(capitalized_words)
```

第一次去女友家，女友叫我带五箱茅台，20条软中华，同时带大红包，怎么办？

第一次去女友家，女友叫你带五箱茅台，20条软中华，同时带大红包，这是一种非常正式和亲近的礼物馈赠方式。以下是一些建议：

- 尊重女友的意愿：虽然这是一种有创意的礼物选择，但在选择礼物之前，最好与女友进行充分的沟通，了解她真正想要的是什么样的礼物。如果女友明确表示不希望收到贵重礼物，那就不要强迫她。
- 根据实际情况做出选择：如果你经济状况较为宽裕，可以考虑选择更加高档的礼物。但如果你的经济能力有限，建议考虑选择适合自己的礼物。例如，可以选择一些小巧、实用且适合女友家庭使用的礼品。
- 注意礼物的包装和交付方式：为了让礼物看起来更加体面和尊重，应该选择适当的包装和交付方式。例如，可以选择礼盒或袋子将礼物包裹好，并在送礼时准备一份礼单，写明礼品名称、数量和价值等信息。同时，还应该及时向女友提供礼物，并感谢她对你的慷慨。
- 感谢女友的关心：女友给你准备这么多礼物，其实也是对你关心和爱护的一种回报。在接受礼物后，可以向女友表达感激之情，并感谢她的关心。

总之，第一次去女友家，选择适当的礼物并保持诚信和礼貌，可以增进彼此间的感情和信任。

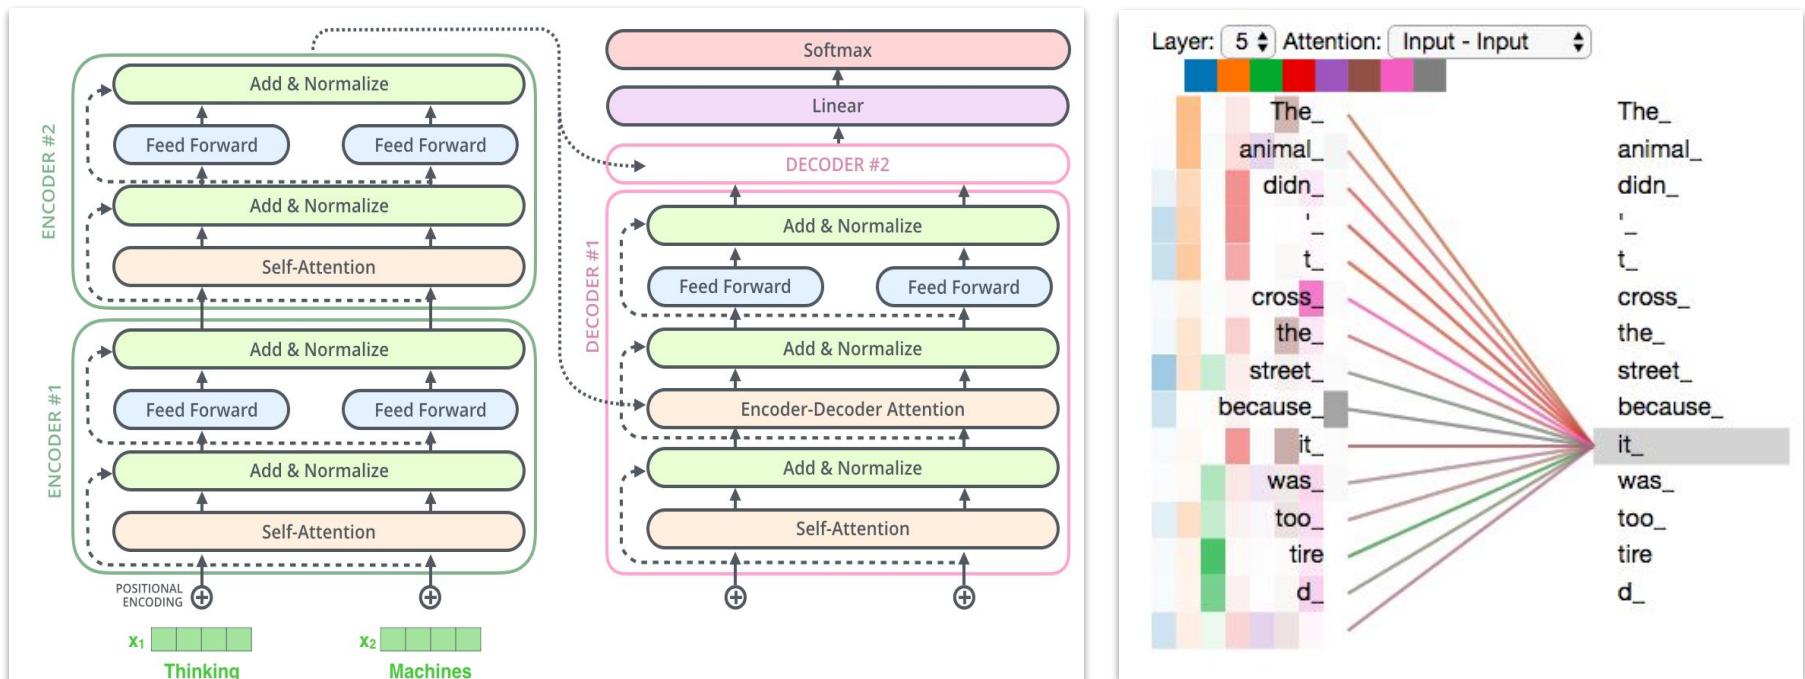
生成式任务简介



语音转图像任务

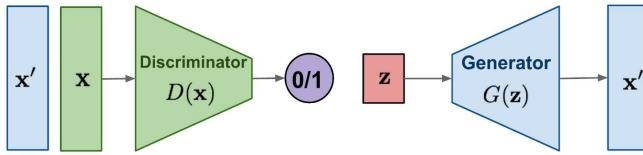
02

生成任务的基础模型——transformer

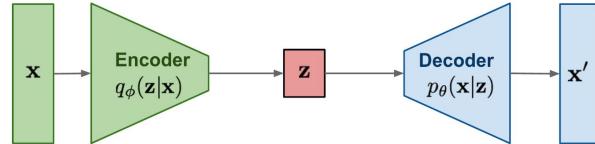


生成任务的基础模型——diffusion model

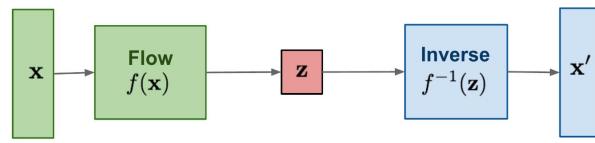
GAN: Adversarial training



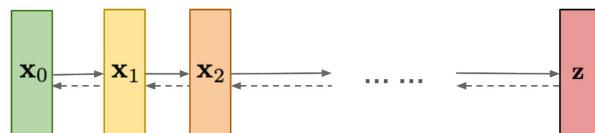
VAE: maximize variational lower bound



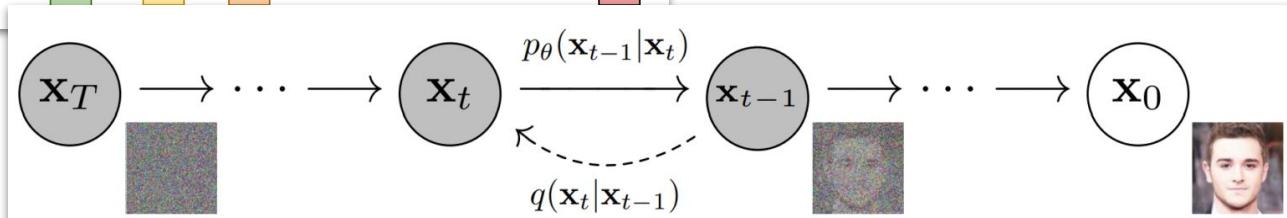
Flow-based models:
Invertible transform of distributions



Diffusion models:
Gradually add Gaussian noise and then reverse



	VAE	Flow	GAN	Diffusion
Pros	Fast Sampling rate. Diverse sample generation	Fast Sampling rate. Diverse sample generation	Fast Sampling rate. High sample generation quality.	High sample generation quality. Diverse sample generation
Cons	Low sample generation quality	Need specialized architecture, low sample generation quality	Unstable training, low sample generation diversity (Mode Collapse)	Low sampling rate



语音转图像第一阶段——语音转文本

- OpenAI whisper模型

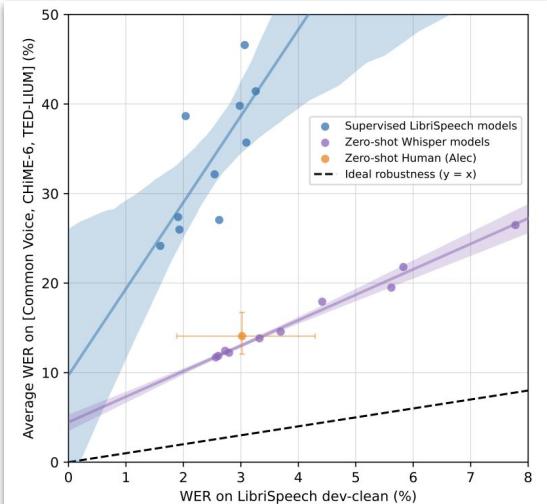
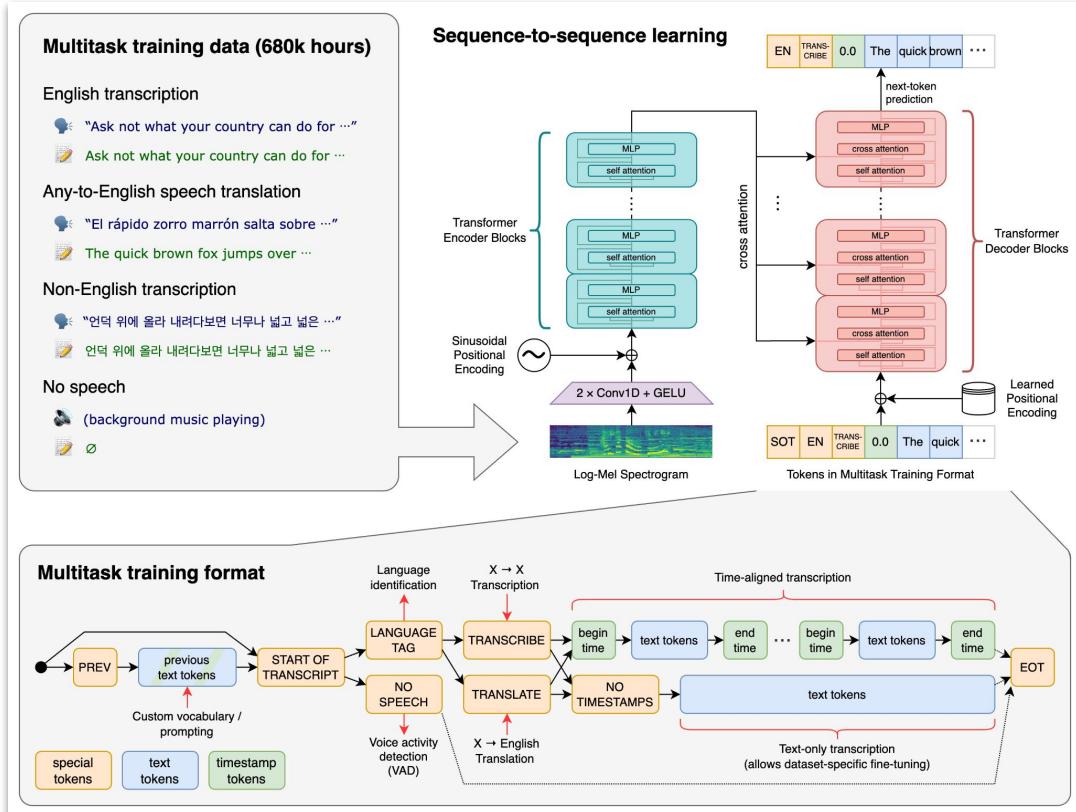
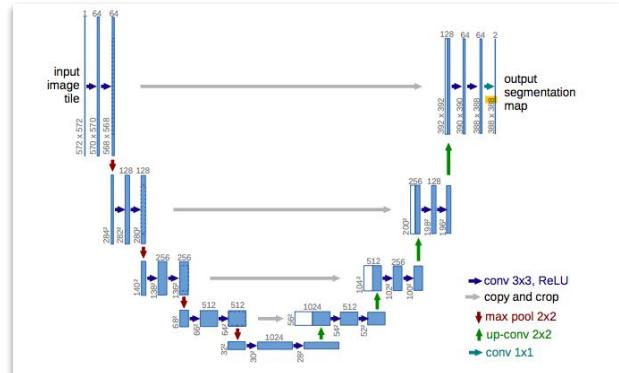
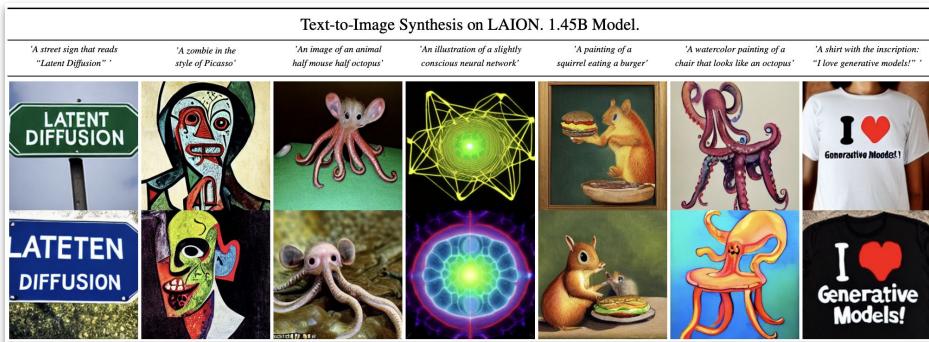
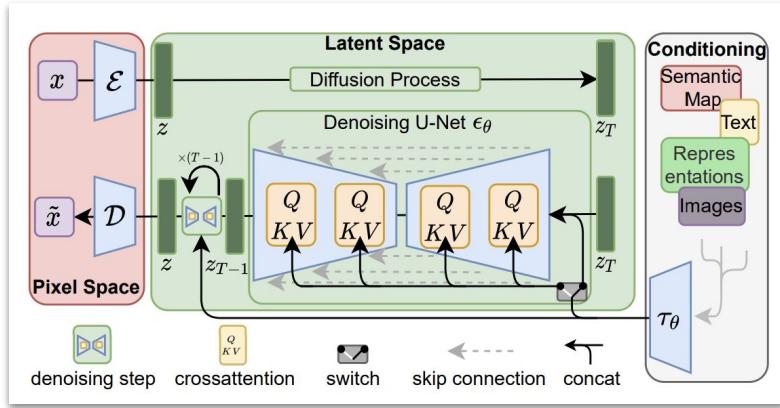


Figure 2. Zero-shot Whisper models close the gap to human robustness. Despite matching or outperforming a human on LibriSpeech dev-clean, supervised LibriSpeech models make roughly twice as many errors as a human on other datasets demonstrating their brittleness and lack of robustness. The estimated robustness frontier of zero-shot Whisper models, however, includes the 95% confidence interval for this particular human.



语音转图像第二阶段——文本转图像

- Stable Diffusion模型
 1. 图像感知压缩(Perceptual Compression)
 2. 潜在扩散模型(Latent Diffusion Models)
 3. 条件机制(Conditional Mechanisms)



使用Jina快速搭建工作流

03



如何用Jina搭建工作流

- 为什么要使用Jina



- Seamless Docker container integration: sharing, exploring, sandboxing, versioning and dependency control via Executor Hub.
- Full observability via OpenTelemetry, Prometheus and Grafana.
- Fast deployment to Kubernetes and Docker Compose.



Ecosystem

- Improved engineering efficiency thanks to the Jina AI ecosystem, so you can focus on innovating with the data applications you build.



Universal

- Build applications that deliver fresh insights from multiple data types such as text, image, audio, video, 3D mesh, PDF with LF's DocArray.

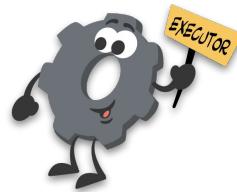
如何用Jina搭建工作流

- Jina的几个核心概念



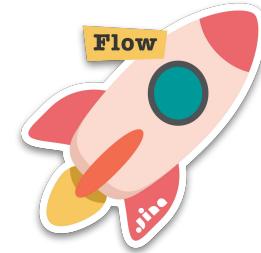
DocArray:

Basic data type to represent a piece of information(text, image, etc)



Executor:

How Jina process Documents.
(Segment, encode, index, etc)



Flow:

How Jina streamlines and distributes Executors

如何用Jina搭建工作流

- 使用 DocArray 表示多模态数据

```
- from dataclasses import dataclass
from typing import List
+ from docarray import dataclass, Document
+ from docarray.typing import Image, JSON, Text

@dataclass
class NewsArticle:
-   content_picture: str
-   title: str
-   paragraphs: List[str]
-   meta: dict
+   content_picture: Image
+   title: Text
+   paragraphs: List[Text]
+   meta: JSON

article = NewsArticle(
    ...
)

+ doc = Document(article)
```

Document: 4065b71e72e10fbafdf60b3e772f84997	
└ Chunks	
└ Document: fba6d9cc388e26d89da09c377998147	
Attribute	Value
parent_id	4065b71e72e10fbafdf60b3e772f84997
granularity	1
tensor	<class 'numpy.ndarray'> in shape (316, 500, 3), dtype: uint8
uri	https://inews.gting.com/newsapp_bt/0/15123866257/1000
modality	image
└ Document: 66f951ac1f33677a190ef67c77251615	
Attribute	Value
parent_id	4065b71e72e10fbafdf60b3e772f84997
granularity	1
text	中国手机品牌在印度何以遭遇“至暗时刻”
modality	text
└ Document: e8d5614b064bf0c19471e63ebe46383d	
Attribute	Value
parent_id	4065b71e72e10fbafdf60b3e772f84997
granularity	1
└ Chunks	
└ Document: 1a0a208bd1421665cd7e894ecc7486bd	
Attribute	Value
parent_id	e8d5614b064bf0c19471e63ebe46383d
granularity	1
text	日前，中国手机品牌荣耀称“由于众所周知的原因撤出印度团队”。近 ...
modality	text
└ Document: 80209132e1a7f2418c1da7752fe11a69	
Attribute	Value

如何用Jina搭建工作流

- 使用DocArray表示多模态数据
- 使用嵌套结构表示不同语义颗粒度的信息
- 支持不同模态的基本操作
- 支持colab和notebook中数据预览
- 支持条件查询
- 支持不同的深度学习框架
- 支持向量表示/预览及不同的向量存储方案

```
r = da.find({'modality': {'$eq': 'D'}})

pprint(r.to_dict(exclude_none=True)) # just for pretty print

[{"id": "92aee5d665d0c4dd34db10d83642aded",
 "modality": "D",
 "tags": {"h": 8.5, "uom": "in", "w": 11.0},
 "text": "paper",
 "weight": 100.0},
 {"id": "1a9d2139b02bc1c7842ecda94b347889",
 "modality": "D",
 "tags": {"h": 22.85, "uom": "cm", "w": 38.0},
 "text": "planner",
 "weight": 75.0}]
```

The screenshot shows a Jupyter Notebook interface with several code cells and their outputs.

- Code Cell 1:** Shows a video player displaying a white rabbit in a grassy field. Below it, the command `from docarray import Document` is shown, followed by code to load a video file and print the tensor shape: `(250, 176, 320, 3)`.
- Code Cell 2:** Shows a video player displaying a white rabbit in a grassy field. Below it, the command `from docarray import Document` is shown, followed by code to load a wav file and display it.
- Code Cell 3:** Shows a video player displaying a white rabbit in a grassy field. Below it, the command `from docarray import Document` is shown, followed by code to load a mp4 file and display it.

如何用Jina搭建工作流

- Executor的内部实现

```
@requests
def transcribe(self, docs: DocumentArray, **kwargs):
    """
    Receives a DocumentArray of Documents containing audio data, either in `Document.uri` or `Document.tensor`. It converts it to text using whisper model and inserts the resulting text in `Document.text`. It is also possible to perform translation of other languages into english. The language can be specified in `Document.tags` as a key-value pair, for instance: `language:en`. List of supported languages: https://github.com/openai/whisper/blob/8cf36f354d4a7a3a2e35403a7a25a442/langs.txt
    :param model_name: the model name used to load whisper. Available model names are: "tiny", "base", "small", "medium", "large".
    """

    for i, doc_ in enumerate(docs):
        model_output = self.model.transcribe(
            doc_.uri if doc_.tensor is None else doc_.tensor, task='translate',
        )
        doc_.text = model_output['text']
        doc_.tags['segments'] = model_output['segments']
        doc_.tags['language'] = model_output['language']
    return docs
```

```
@requests(on='/')
def generate(self, docs: DocumentArray, parameters: dict, **kwargs):
    """
    Generate images from text using the StableDiffusion model.
    This endpoint expects documents with text prompt stored in `Document.text`, will generate up to `num_images` images represented as Document objects. The Document objects contain the image data stored in the `Document.uri` field and are added to the query `Document.matches` field.
    The `num_images` parameter can be specified as a key-value pair in `parameters` of the request.
    """

    num_images = parameters.get('num_images', 1)

    for document in docs:
        self.generate_image_from_document(document, num_images)

    return docs

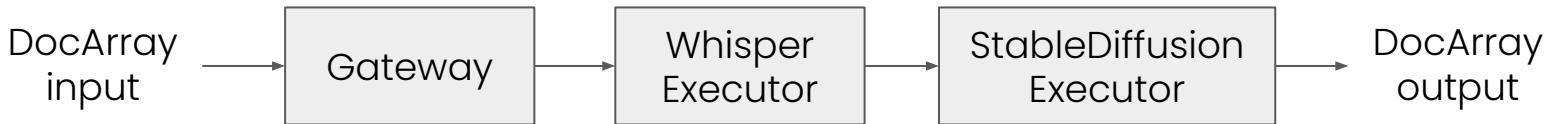
def generate_image_from_document(
    self, document: Document, num_images: int
) -> Document:
    with torch.autocast('cuda'):
        generated_imgs = self.diffusion([document.text] * int(num_images)).images

    for img in generated_imgs:
        _generated_document = Document(
            tags={
                'text': document.text,
                'generator': self.__class__.__name__,
            }
        ).load_pil_image_to_datauri(img)

    document.matches.append(_generated_document)
```

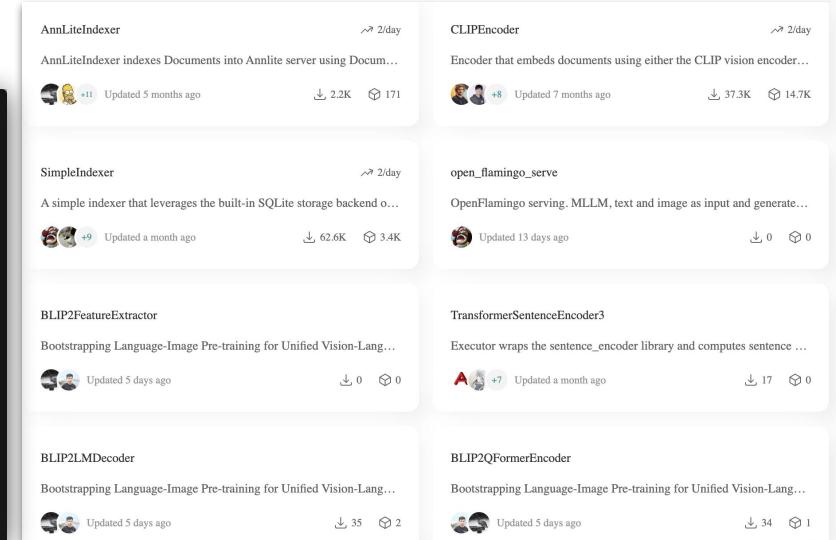
如何用Jina搭建工作流

- Flow的构成



- Flow yaml

```
jtype: Flow
executors:
  - name: whisper
    uses: jinahub+docker://WhisperExecutor
    uses_with:
      model_name: medium
      gpus: all
      jcloud:
        resources:
          gpu: 1
          memory: 2G
  - name: diffusion
    uses: jinahub+docker://StableDiffusionExecutor
    uses_with:
      auth_token: ${{ ENV.HF_TOKEN }}
      timeout_ready: -1 # slow download speed often leads to timeout
      replicas: 2
      jcloud:
        resources:
          gpu: 1
          memory: 16G
```



如何用 Jina 搭建工作流

- 启动 Flow 并发送请求

```
● ● ●  
import os  
from jina import Flow  
from executors.whisper import WhisperExecutor  
from executors.stablediffusion import StableDiffusionExecutor  
  
hf_token = os.getenv('HF_TOKEN')  
  
f = (  
    Flow(port=54322)  
    .add(uses=WhisperExecutor, timeout_ready=-1, uses_with={'model_name': 'large'})  
    .add(uses=StableDiffusionExecutor, uses_with={'auth_token': hf_token})  
)  
if __name__ == '__main__':  
    with f:  
        f.block()
```



```
from jina import Client  
from docarray import Document  
client = Client(host='localhost:54322')  
docs = client.post('/', inputs=[Document(uri='audio.wav') for _ in range(1)])  
for img in docs[0].matches:  
    img.load_uri_to_image_tensor()  
  
docs[0].matches.plot_image_sprites()
```

Demo

04

jina

Berlin · San Jose · Beijing · Shenzhen

大家：問題那麼多有沒有禮物啊

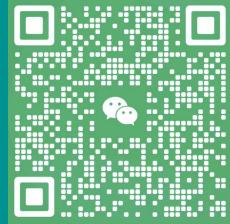


只有两道题



而且有周边

领周边指路 ➔



扫一扫上面的二维码图案，加我为朋友。

