

Production LLM Deployment

DS Pipelines in Minutes: An Example...

PDF Summary Pipeline

Using Seaplane and GPT3.5



Don't believe us, believe them!

"A straightforward platform for data science deployments without getting caught up in configuration tasks. A data scientist's dream 😊"

Kjeld Oostra AI Engineer @ Entropical

What is Seaplane

A powerful Python SDK to take your data science experiments to production in minutes without having to wrangle any cloud infrastructure.

Seaplane is jam packed with functionality to enable data scientist to do their best work. Such as built-in SQL, Vector-Store, Object Store, Large models and much more!




PDF Summary Application

- An API entry point available at [/demo-input](#)
- A pre-processing step to download, extract the text from our PDF and create the prompt for GPT-3.5
- An inference step to create the summary using GPT-3.5
- A database step to store the result in Seaplane's managed global SQL database



Directed Acyclic Graph



How to build an app on Seaplane

- Applications `@app` - HTTP, Kafka, Cron job entry points
- Tasks `@task` - Individually scalable containers
- Task Resources
 - SQL
 - Vector-DB
 - Object Store
 - LLMs
 - Much more!

Let's Start Building



The App

- Creates the HTTP endpoint
- Creates the DAG i.e the flow of the data
- Starts the application

```
1 from seaplane import config, app, start
2
3 api_keys = {
4     "SEAPLANE_API_KEY": os.getenv("SEAPLANE_API_KEY"),
5     "OpenAI_API_KEY": os.getenv("OpenAI_API_KEY"),
6 }
7
8 config.set_api_keys(api_keys)
9
10 @app(path='/demo-input', method='POST', id='pdf-summary')
11 def my_smartpipe(body):
12     # wire the tasks together in a DAG
13     prompt = pre_processing(body),
14     summary = inferencing(prompt),
15     database(summary)
16
17 start()
```

Preprocessing @task

- Downloads PDF
- Extract Text
- Create Prompt

```
1 from seaplane import task
2 import requests
3 from PyPDF2 import PdfReader
4
5 @task(type='compute', id='pre-processing')
6 def pre_processing(data):
7     # get the URL from the request
8     url = data['url']
9
10    # download PDF and extract text here
11
12    # construct prompt
13    prompt = "write a summary of the following text. Make
14             sure to maintain the scientific tone in the paper: "
15             + pdf_text
16
17    # pass the required information to the next step
18    return({
19        'url' : url,
20        'prompt' : prompt
21    })
```


Inference @task

- Runs the inference
- OpenAI and other LLMs available through the SDK
- Local models such as MPT-30B, Bloom, Bloomz - Shared resource hosted by Seaplane.

```
1 from seaplane import task
2
3 @task(type='inference', id='pdf-inferencer', model="gpt-3.5")
4 def inferencing(data, model):
5     prompt = data['prompt']
6     url = data['url']
7
8     # construct the model parameters including the prompt from
9     # the prev step
10    params = {
11        "model": "gpt-3.5-turbo",
12        "messages": [{"role": "user", "content": prompt}],
13        "temperature": 0.7,
14    }
15
16    # run the inference request
17    result = model(params)
18
19    # return the inferred result plus input and parameters
20    result['url'] = url
21    result['prompt'] = prompt
22    return(str(result).encode())
```

Database @task

- Stores the result in Seaplane SQL
- DB connection managed by SDK

```
1 from seaplane import task, sql
2
3 sql_access = {
4     "username": "<YOUR-USERNAME>",
5     "password": "<YOUR-PASSWORD>",
6     "database": "<YOUR-DB-NAME>",
7 }
8
9 @task(type='compute', id='pdf-summary-db', sql=sql_access)
10 def database(data, sql_db):
11     # insert into table
12     sql_db.insert('' INSERT INTO pdf_summaries
13                   (url, prompt, choices)
14                   VALUES
15                   (%s,%s,%s)
16                   ''', [
17                   data["url"], data['prompt'], data['summary']]
18     )
```

Deploying

- Set up API gateway
- Deployed 3 containers
- Wired them all together
- Set up and configured database
- Scalable serverless multi-cloud infrastructure

```
$ seaplane deploy

[Seaplane]

    Seaplane Apps version 0.3.69

[Seaplane] 🚀 Assign Task pre-processing to App: pdf-summary
[Seaplane] 🚀 Assign Task pdf-inferencer to App: pdf-summary
[Seaplane] 🚀 Assign Task pdf-summary-db to App: pdf-summary
[Seaplane] Apps build successfully!

[Seaplane] Requesting access token...
[Seaplane] Requesting access token...
[Seaplane] Deploy for task pre-processing done
[Seaplane] Deploy for task pdf-inferencer done
[Seaplane] Deploy for task pdf-summary-db done
[Seaplane] 🏁 Deployment complet
```

```
~/Desktop/pdf-summary 04.8340
ls
build          pdf-summary    poetry.lock    pyproject.toml test.ipynb

~/Desktop/pdf-summary
seaplane deploy
[Configuration file exists at /Users/robbedeker/Library/Application Support/py Poetry, reading this directory. Symbolic
links are moving from configuration files to /Users/robbedeker/Library/Preferences/py Poetry, as support for the legacy director
y will be removed in an upcoming release.]
[Seaplane]

      Seaplane App version 0.3.00

[Seaplane] # Assign Task pre-processing to App: pdf-summary
[Seaplane] # Assign Task pdf-interlacer to App: pdf-summary
[Seaplane] # Assign Task pdf-summary-db to App: pdf-summary
[Seaplane] App build successfully!

[Seaplane] Requesting access token...
[Seaplane] Requesting access token...
[Seaplane] Deploy for Task: pre-processing done
[Seaplane] Deploy for Task: pdf-interlacer done
[Seaplane] Deploy for Task: pdf-summary-db done
```

<https://www.youtube.com/watch?v=M4nhkC1E2Uw>



Sign up for the beta today and we will double your credits*



seaplane.io/join-the-seaplane-beta

*some restrictions apply