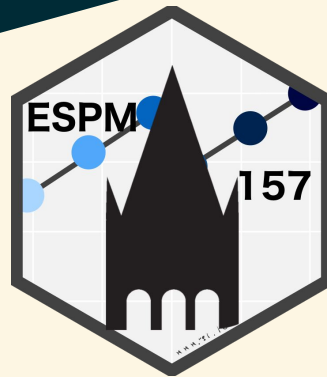


Jupyter-AI & Open LLMs

in large science classes



Carl Boettiger
UC Berkeley



- 122 students
- ~ 2/3 bio sciences
- 1/3 data science
- Active learning classroom



ESPM-157: Data Science for Global Change Ecology

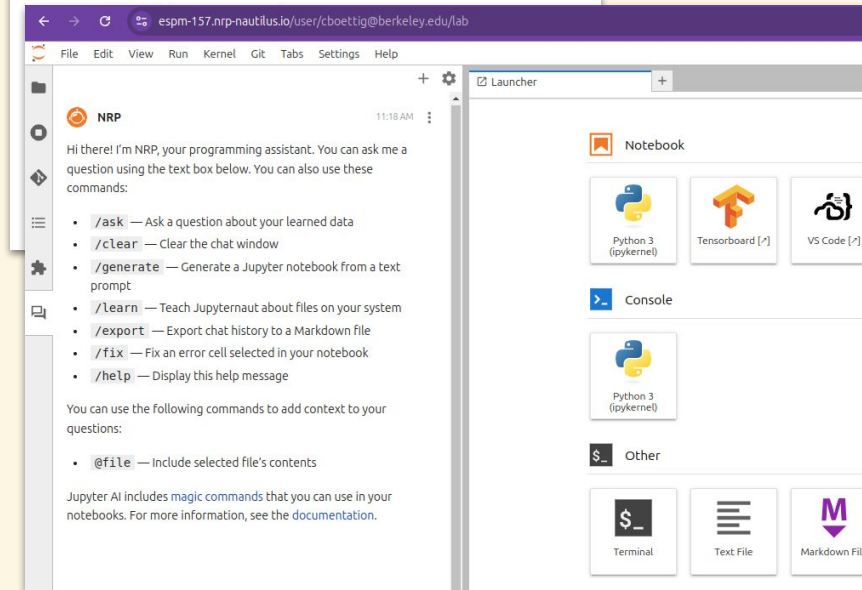
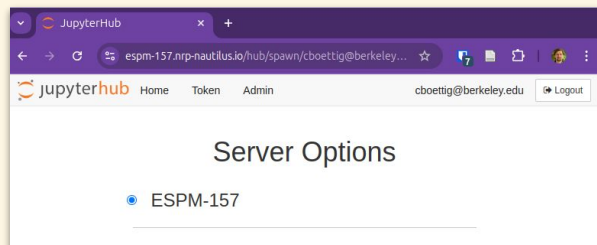
Motivations

LLMs are here. Am I preparing my students for this future?

- Can we go further, faster by **coding** with LLM assistants?
- What are **environmental** implications of LLMs?
- What are the **ethical** implications?
- Can we make LLM use **safer** & more **reliable**?

Computing Resources

[nvidia.com/gpu](https://www.nvidia.com/gpu): "1" x 122 students!



```
! jupyter-ai.yml
1  name: base
2  channels:
3    - conda-forge
4  dependencies:
5    - jupyterlab-myst
6    - jupyter-ai
7    - jupyter-resource-usage
8    - jupyter-vscode-proxy
9    - pypdf
10   - langchain
11   - streamlit
12   - ibis-duckdb
13   - leafmap
14   - odc-stac
15   - pystac-client
16   - planetary-computer
17   - mystmd
18   - pip
19   - pip:
20     - jupyter-tensorboard-proxy
21     - git+https://gitlab.nrp-nautilus.io
22     - langchain-ollama
23     - langchain-openai
24     - langchain-mistralai
25     - git+https://github.com/eodaGmbH/py
26
```

Environment.yml
Docker image

```
! values.yaml
1  hub:
2    db:
3      type: sqlite-pvc
34     pvc:
35       accessModes:
36         - ReadWriteOnce
37       storage: 10Gi
38       storageClassName: rook-ceph-block
39 resources:
40   limits:
41     cpu: "2"
42     memory: 4Gi
43   requests:
44     cpu: 100m
45     memory: 512Mi
46 networkPolicy:
47   enabled: false
48 service:
49   type: ClusterIP
50 chp:
51   resources:
52     limits:
53       cpu: "6"
54       memory: 10Gi
55       nvidia.com/gpu: "1"
56     requests:
57       cpu: "1"
58       memory: 512Mi
59       nvidia.com/gpu: "1"
singleuser:
```

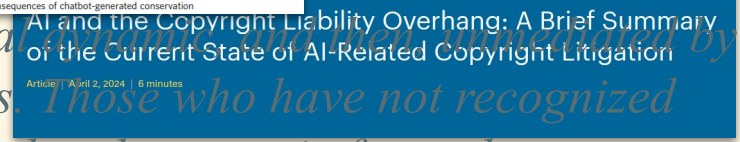
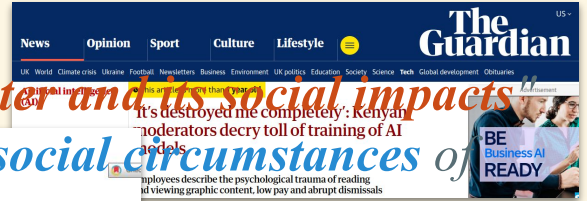
k8s Helm chart

Challenges with respect to AI

- those who focus practically on such things as "the computer and its social impacts" but who fail to **look behind technical things** to notice the **social circumstances** of their development, deployment, and use.
- Energy footprint?
- Ethical concerns:

*This view provides an **and** **note** to naive **technological determinism** – the idea that technology develops as the sole result of an internal logic, unaffected by any other influence, molds society to fit its patterns. Those who have not recognized the ways in which technologies are shaped by social and economic forces have not gotten very far.*

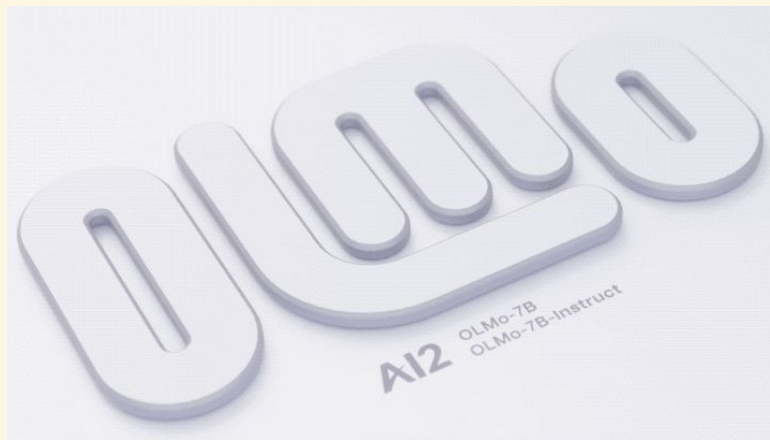
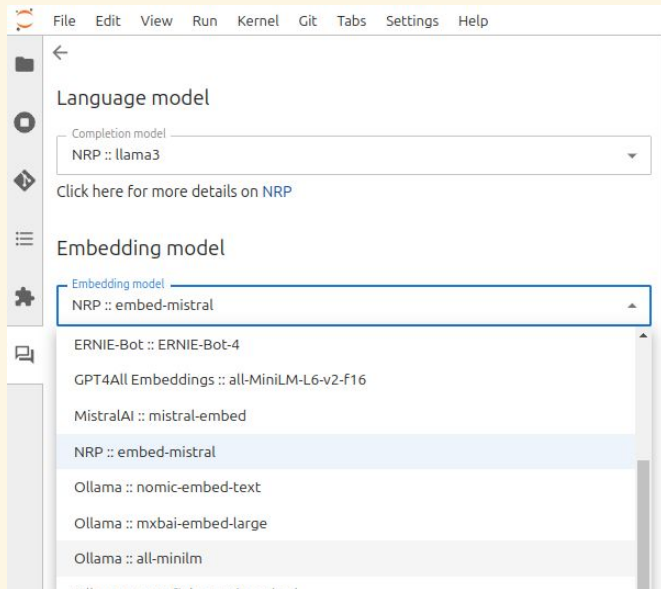
- Langdon Winner, 1980



Local, open models:

```
$ ollama serve  
$ ollama pull nomic-embed-text
```

CO2 footprint data, not
hype:



Researcher developed + hosted

```
! llm.yml
6 spec:
15   template:
20     spec:
21       affinity:
30         containers:
31           - args:
32             - -m
33               - vllm.entrypoints.openai.api_server
34             - --port
35             - "5000"
36             - --host
37             - 0.0.0.0
38             - --download-dir
39             - /workspace/.cache/huggingface/hub
40             - --model
41             - gorilla-llm/gorilla-openfunctions-v2
42             - --tensor-parallel-size
43             - "2"
44             - --trust-remote-code
45             - --enable-auto-tool-choice
46             - --tool-call-parser
47             - llama3_json
48 image: vllm/vllm-openai:v0.6.3
49 imagePullPolicy: IfNotPresent
50 name: gorilla-openfunctions-v2
51 command:
52   - python3
53 resources:
54   limits:
55     cpu: "5"
56     memory: 36Gi
57     nvidia.com/rtxa6000: "2"
58 requests:
59   cpu: "1"
60   memory: 36Gi
61   nvidia.com/rtxa6000: "2"
62
```

The screenshot shows the LiteLLM Model Hub interface. On the left, there is a sidebar with navigation options: Virtual Keys, Test Key, Usage, API Reference, and Model Hub (which is highlighted). The main area is titled "Model Hub" and includes a "Make Public" button. It displays a grid of eight model cards, each with a title, a "View more" link, and a list of capabilities and constraints.

Model Name	Mode	Supports Function Calling	Supports Vision	Max Input Tokens	Max Output Tokens
ph13		No	Yes	N/A	N/A
mixtral		No	No	N/A	N/A
llava		No	Yes	N/A	N/A
embed-mistral	embedding	No	No	N/A	N/A
gemma2		No	Yes	N/A	N/A
llama3-sdsc		No	Yes	N/A	N/A
gorilla	chat	Yes	No	N/A	N/A
llama3		No	Yes	N/A	N/A

> base_url = "https://llm.nrp-nautilus.io/"

Challenges with respect to AI

- LLMs don't

The image shows a chat interface on the left and a terminal window on the right. The chat shows a user named 'jovyan' asking for Python code to plot, and a response from 'NRP' providing a detailed Python script. The terminal window shows the execution of this script, displaying the output of the metadata and a 3D map visualization of a city with buildings.

Chat Conversation:

jovyan 7:19 PM
/ask provide example code with pmtiles

NRP 7:20 PM
To create a 3D map of buildings using PMTiles, you can use the leafmap library, which provides an interface to visualize PMTiles in 3D. Here's an example code that demonstrates this process:

```
import leafmap.maplibregl as leafmap

# URL of the PMTiles file
url = "https://data.source.coop/cholmes/overture/ove

# Get the metadata of the PMTiles file
metadata = leafmap.pmtiles_metadata(url)
print(f"layer names: {metadata['layer_names']}")
print(f"bounds: {metadata['bounds']}")

# Create a map object
m = leafmap.Map(
    center=[-74.0095, 40.7046], zoom=16, pitch=60, b

)

# Add a basemap
m.add_basemap("OpenStreetMap.Mapnik")

# Define a style for the 3D buildings
style = {
    "layers": [
        {
```

Terminal Output:

```
m.add_pmtiles(
    url,
    style=style,
    visible=True,
    opacity=0.8,
    tooltip=True,
)

# Display the map
m

layer names: ['buildings']
bounds: [-179.999966, -83.6500138, 179.9972908, 82.525

[1]:
```

The terminal output shows a 3D map visualization of a city with buildings. The map is rendered in a perspective view, showing a dense urban area with a river and a bridge. The buildings are colored in shades of blue and yellow, and the terrain is green. The map is displayed in a window titled 'Terminal 1'.

/learn docs
/ask ...


```
st.title("SQL demo")
```

```
## dockerized streamlit app wants to read from os.
```

```
api_key = os.getenv("LITELLM_KEY")
```

```
if api_key is None:
```

```
    api_key = st.secrets["LITELLM_KEY"]
```

```
parquet = st.text_input("parquet file", "https://
```

```
eng = sqlalchemy.create_engine("duckdb:///memory:
```

```
con = ibis.duckdb.from_connection(eng.raw_connecti
```

```
tbl = con.read_parquet(parquet, "mydata")
```

```
# langchain can also talk to this connection and s
```

```
db = SQLiteDatabase(eng, view_support=True)
```

```
# Build the template for system prompt
```

```
template = '''
```

```
You are a {dialect} expert. Given an input questio  
Always return all columns from a query (select *)
```

```
Wrap each column name in double quotes (") to deno  
Pay attention to use only the column names you can
```

```
Be careful to not query for columns that do not ex  
Also, pay attention to which column is in which ta
```

```
Pay attention to use today() function to get the c  
Respond with only the SQL query to run. Do not re
```

```
Only use the following tables:
```

```
{table_info}
```

```
Question: {input}
```

```
'''
```

```
from langchain_community.llms import OpenAI
```

RAG & TAG – Building LLM Agents

The screenshot shows a web interface for a Streamlit application. At the top, it says 'Spaces | cboettig/streamlit-demo' with a 'like 0' button and a 'Running' indicator. The main title is 'SQL demo'. Below the title, there is a 'parquet file:' label and a text input field containing the URL 'https://espm-157-f24.github.io/spatial-carl-amanda-tyler/new_haven_stats.parquet'. Below the input field, there is a code editor showing a SQL query: 'SELECT grade, AVG(ndvi) AS mean_ndvi FROM mydata GROUP BY grade;'. To the right of the code editor, there is a table with the following data:

	grade	mean_ndvi
0	A	0.7308
1	C	0.5472
2	D	0.4866
3	B	0.6093

At the bottom of the interface, there is a text input field containing the question 'What is the mean ndvi by grade?' and a right-pointing arrow button.

Project Contributions & Future engagement

- Teach **safe LLM** use, not **AI abstinence**
- **Open models** mitigate some risks (energy, privacy)
- NRP as platform for collaborative innovation!
 - **Shared** hardware, software, models