

## **EE/CprE/SE 492 WEEKLY REPORT 2**

**Sept 6 - Sept 19**

**Group number: 20**

**Project title: GridGPT**

**Client &/Advisor: Gelli Ravikumar**

**Team Members/Role:**

- **Tin Ngo -> AI Integration Specialist**
- **Jackson Phillips -> AI Integration Specialist**
- **Emma Heithoff -> Power Systems Specialist**
- **Eddy Andrade -> Frontend Lead**
- **Nick Doty -> Power Systems Specialist**

### ○ **Weekly Summary**

Over the past two weeks, our team made significant progress on various aspects of the GridGPT project. Tin focused on integrating OpenAI API functionalities into our system by developing modular classes and setting up a Flask API to generate text outputs, while also looking into HPC access. Jackson researched training Llama 3 and found a Colab notebook suitable for use with the university's HPC cluster. Emma laid the foundation for running power system simulations and prepared a plan for dataset creation in parallel to AI training. Eddy began familiarizing himself with Git/GitLab and started preparing for frontend development after VM cloning was completed. Nick focused on setting up his VM environment and reviewing DSS research data to align with Emma's simulation tasks.

### ○ **Past week accomplishments**

Tin Ngo: I got more research done on OpenAI and how to use the API. I implemented two classes because I want to make it as modular as possible early on so we can reuse it. The first class that I created is called APIAuthenticator. This holds the API Key using the environment variables. This will be the parent class for many other courses to have authentication. The second class I made

was OpenAI API. You can instantiate this class with a model that you specify. It has a `generate_text` method that allows you to input a prompt and receive a streamed output. To use these classes, I made a Flask API endpoint that creates the OpenAI API object and generates text very similarly to ChatGPT. All you have to do is provide the prompt and model in JSON format. I added a linting step to my Flask app in the CI/CD pipeline to follow the project's conventions. I then looked into accessing HPC. I also worked on creating weekly deliverables as my advisor has asked for them.

[https://git.ece.iastate.edu/sd/sddec24-20/-/merge\\_requests/6/diffs#4c71286b4bc03f9201fffc2e49826ea574f35cdf](https://git.ece.iastate.edu/sd/sddec24-20/-/merge_requests/6/diffs#4c71286b4bc03f9201fffc2e49826ea574f35cdf)

```
(venv) tinngo@tinngo sddec24-20 % curl -X POST http://127.0.0.1:5000/generate_text -H "Content-Type: application/json" -d '{"prompt": "Say this is a test", "model": "gpt-3.5-turbo-0125"}'
```

```
data: This
```

```
data: is
```

```
data: a
```

```
data: test
```

```
data: .
```

```
Also note that slurm accounts are being reorganized to support better resource-management and improve usage reporting. You may view your available slurm accounts by running 'slurm-account-selector'. This will also allow you to select the default account to be used when submitting jobs without explicitly using -A/--account.
```

```
#####
```



```
The next Nova maint is scheduled for Monday Jan 6th, 2025. We plan on moving equipment around the data center so this may be more than 1 day.
```

```
#####
```

```
Updating slurm account cache  
[tinngo56@nova ~]$ cd work/gelli  
-bash: cd: work/gelli: No such file or directory  
[tinngo56@nova ~]$ cd /work/gelli  
[tinngo56@nova gelli]$ ls  
gelli kkumar sddec24-20 tinngo56  
[tinngo56@nova gelli]$
```

<b>Open</b> 5	Closed 0	All 5	Filter by milestone name	Due soon ▾	New milestone
<hr/>					
<b>Sprint 1</b>			<div style="border: 1px solid green; width: 100%; height: 10px;"></div>		⋮
Aug 21, 2024–Sep 19, 2024			1 Issue · 1 Merge request	100% complete	
<span>Open</span> sd / sddec24-20					
<hr/>					
<b>Sprint 2</b>			<div style="border: 1px solid gray; width: 100%; height: 10px;"></div>		⋮
Sep 19, 2024–Oct 3, 2024			3 Issues · 0 Merge requests	0% complete	
<span>Upcoming</span> sd / sddec24-20					
<hr/>					
<b>Sprint 3</b>			<div style="border: 1px solid gray; width: 100%; height: 10px;"></div>		⋮
Oct 3, 2024–Oct 17, 2024			3 Issues · 0 Merge requests	0% complete	
<span>Upcoming</span> sd / sddec24-20					
<hr/>					
<b>Sprint 4</b>			<div style="border: 1px solid gray; width: 100%; height: 10px;"></div>		⋮
Oct 17, 2024–Oct 31, 2024			2 Issues · 0 Merge requests	0% complete	
<span>Upcoming</span> sd / sddec24-20					
<hr/>					
<b>Sprint 5</b>			<div style="border: 1px solid gray; width: 100%; height: 10px;"></div>		⋮
Oct 31, 2024–Nov 14, 2024			0 Issues · 0 Merge requests	0% complete	
<span>Upcoming</span> sd / sddec24-20					

The screenshot shows a task board interface with a header bar containing a dropdown menu with 'Tin', a notification icon with '15', and a settings icon. The tasks are listed in a vertical stack, each with a title, an ID, and a sprint assignment. The tasks are:

- #53
- Create Database to save messages overtime**  
#52
- Use data from database to use as a prompt for AI**  
#51 ◇ Sprint 4
- Implement feedback changes**  
#50 ◇ Sprint 3
- Write up documentation for front end. Meet with Group and create additional work based on feedback**  
#49 ◇ Sprint 3
- Host small LLM from HuggingFace onto HPC**  
#48 ◇ Sprint 2 AI 
- Create Weekly Deliverables**  
#47 

Jackson Phillips: I spent time this week looking into how to train Llama 3 specifically. I found a google colab notebook that goes in depth on llama 3 specifically and doesnt require any extra UI, programs, or websites which I think we will need because we need to train it using the university's HPC cluster and not on a normal pc.

## ▼ Train the model

Now let's use Huggingface TRL's SFTTrainer! More docs here: [TRL SFT docs](#). We do 60 steps to speed things up, but you can set `num_train_epochs=1` for a full run, and turn off `max_steps=None`. We also support TRL's DPOTrainer!

```
[ ] from trl import SFTTrainer
    from transformers import TrainingArguments
    from unsloth import is_bfloat16_supported

    trainer = SFTTrainer(
        model = model,
        tokenizer = tokenizer,
        train_dataset = dataset,
        dataset_text_field = "text",
        max_seq_length = max_seq_length,
        dataset_num_proc = 2,
        packing = False, # Can make training 5x faster for short sequences.
        args = TrainingArguments(
            per_device_train_batch_size = 2,
            gradient_accumulation_steps = 4,
            warmup_steps = 5,
            # num_train_epochs = 1, # Set this for 1 full training run.
            max_steps = 60,
            learning_rate = 2e-4,
            fp16 = not is_bfloat16_supported(),
            bf16 = is_bfloat16_supported(),
            logging_steps = 1,
            optim = "adamw_8bit",
            weight_decay = 0.01,
            lr_scheduler_type = "linear",
            seed = 3407,
            output_dir = "outputs",
        ),
    )
```

Emma Heithoff: I researched a plan for the dataset to be created in parallel to the AI specialists training the model, and the following are the general steps. First, we will run OpenDSS simulation using AltDSS-Python. We will extract what we need from results (voltages, currents, line losses, other electrical values etc.) and place into csv form. The AI model will need a normalized dataset method and a training accuracy metric. I plan to track my milestones on Git much better so next week's milestones will be added on Friday. I have simulation code using Alt-DSS planned for analysis capabilities for our GridPilot use case.. I will implement these simulations after meeting with Nick and Jackson/Tin, but, this week's work set a foundation I was previously lacking.

Eddy Andrade: We received a notification from our client that the VMs have been cloned, which allows me to begin working on the project. We as a group decided on who was working on what VM. To get started, I needed to become familiar with working with Git and GitLab, especially with branching out of the main branch of the project. This is to avoid any major issues that could impact the project down the line. It is vital for me to regain my understanding of branching, merging, and

running CI/CD on the project as I progress through and continue working on Frontend.

Nick Doty: After we were each assigned our respective VMs, I have begun to get familiar with how it operates as it has been awhile since I have had to use a VM. I have also refreshed myself with our DSS research data

○ **Pending issues**

Tin Ngo: No Pending issues

Jackson Phillips: No issues.

Emma Heithoff: No issues.

Eddy Andrade: Not sure if this may become an issue, but when I branched off on my test branch the CI/CD pipeline failed at the security part of the pipeline. Not sure if when I branched from the HEAD, it was from a version where the security was working or not but something worth mentioning.

Nick Doty: N/A

○ **Individual contributions**

<b><u>NAME</u></b>	<b><u>Individual Contributions</u></b> <i>(Quick list of contributions. This should be short.)</i>	<b><u>Hours this week</u></b>	<b><u>HOURS cumulative</u></b>
Tin Ngo	Developed classes & app to use openAI API, explored HPC, created weekly deliverables	22	101
Jackson Phillips		6	75
Emma Heithoff	Reset foundation for running simulations and staying connected with all parts of project	8	69
Eddy Andrade	Familiarizing with Git/GitLab; Reading up Next.js documentation	6	72

Nick Doty	VM setup an DSS research refresher	6	72
-----------	------------------------------------	---	----

- **Plans for the upcoming week**

Tin Ngo: Work on hosting a small LLM onto HPC, and creating the classes and application to use the LLM.

Jackson Phillips: I want to look further into HPC and how we run things on it.

Emma Heithoff: With running simulations to save power system data, I need to meet a few times or quickly touch base. I will talk to Nick to detail our role's timeline together then run simulations using Alt-DSS. I will talk to Professor Gelli or the graduate students either before running simulations or with questions from simulating to clarify dataset creation questions. Also, I will talk to Jackson & Tin on our collaboration of adding simulations to a dataset and their timeline for model training. I plan to check in with Nick and Jackson/Tin often this semester by meeting or quick messages.

Eddy Andrade: Read up on documentation with Next.js to begin working on implementing the chat window. Main priority is to figure out how to add a window alongside a page.

Nick Doty: I plan to meet with Emma to flesh out our semester timeline and additionally meet up with Tin and Jackson with Emma to coordinate the input of DSS files into the AI model.

- **Summary of weekly advisor meeting**

We were not able to meet in person with Professor Gelli due to scheduling conflicts. To substitute, we each emailed our weekly work, detailed explanation of our progress, questions, and tasks for the upcoming week.