



National
College of
Ireland

Event-Driven Smart Agriculture Monitoring Using AWS Cloud Services

MSc Research Project
MSc in Cloud Computing

VIVEKANANDA KANUKULA
Student ID: X23340312

School of Computing
National College of Ireland

Supervisor: **Jorge Mario Cortes Mendoza**

Student Name:	VIVEKANANDA KANUKULA
Student ID:	X23340312
Programme:	MSc in Cloud Computing
Year:	2025
Module:	MSc Research Project
Supervisor:	Jorge Mario Cortes Mendoza
Submission Due Date:	11/08/2025
Project Title:	Event-Driven Smart Agriculture Monitoring Using AWS Cloud Services
Word Count:	1303
Page Count:	13

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	vivekananda
Date:	11th August 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	<input type="checkbox"/>
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	<input type="checkbox"/>
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	<input type="checkbox"/>

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Event-Driven Smart Agriculture Monitoring Using AWS Cloud Services

VIVEKANANDA KANUKULA
X23340312

1 Introduction

This part describes the software, hardware, and cloud service-based needs of the implementation of the Event-Driven Smart Agriculture Monitoring Using AWS Cloud Services system. Its arrangement will provide a seamless mix of IoT devices, machine learning forecasts, and real-time dashboards with the help of AWS managed services and Python-based developer tools. It also lists the environment variables and python libraries needed to install, deploy, and sustain the project.

2 System Configuration

2.1 Software Requirements

- python : 3.10+
- Dashboard frame work: Streamlit v1.33+
- Development Environment: AWS Cloudshell, Jupyter Notebook (SageMaker Studio)
- MQTT Library:paho-mqtt v1.6+
- GitHub Actions
- Data Processing Libraries:pandas v2.2+, numpy v1.26+
- Learning Library Machine: scikit-learn v1.4+
- AWS CLI / SDK AWS CLI v2, boto3 v1.34+

2.2 Hardware/Hosting

- Local: laptop with \geq 8GB RAM
- EC2 Instance (Dashboard): t3.micro or t3.small

3 AWS Services and Roles

The following AWS services were used in the project, along with their roles and key configurations:

- AWS IoT Core — Secure MQTT message broker to ingest sensor data. Key configurations: Create Thing, certificates, and policy; IoT Rule to trigger Lambda.
- AWS Lambda — Serverless processing that integrates IoT, DynamoDB, S3, SageMaker, and SNS. Key configurations: Python 3.10 runtime; least-privilege IAM role; environment variables set.
- Amazon SageMaker — Hosts the ML model (RandomForestClassifier) to predict irrigation. Key configurations: Model file (.joblib) and inference script (predict.py); endpoint instance type ml.t3.medium.
- Amazon DynamoDB — Stores recent sensor data and predictions for real-time access. Key configurations: Partition key device_id; sort key timestamp.
- Amazon S3 — Archives historical sensor and prediction data. Key configurations: Lifecycle policy to migrate older records to lower-cost storage.
- Amazon SNS — Sends SMS notifications when irrigation is needed. Key configurations: Create topic; phone number subscriptions.
- Amazon EC2 — Hosts the real-time Streamlit dashboard. Key configurations: Amazon Linux; security group open port 8501.

4 Local Environment Setup

1. AWS Account Create and Sign In

Registered an aws account with <https://aws.amazon.com> with an email.

2. Network setup Multi-Factor Authentication (MFA) log-on safeguarding.

3. Log in to aws management console.

4. Access CloudShell AWS

5. Launched AWS CloudShell console by going through AWS Console to execute AWS Command line Interface without having to install any locally.

confirmed CLI version of AWS `aws -v`.

6. Install Configure Dependencies on CloudShell

7. Installed python libraries (boto3, paho-mqtt, pandas, numpy, plotly, streamlit).

8. Set access keys on AWS CLI with aws configure.

9. AWS IoT Core Set up

agriculture-IoT-Room device was a new IoT Thing (virtual device) created in AWS IoT Core.

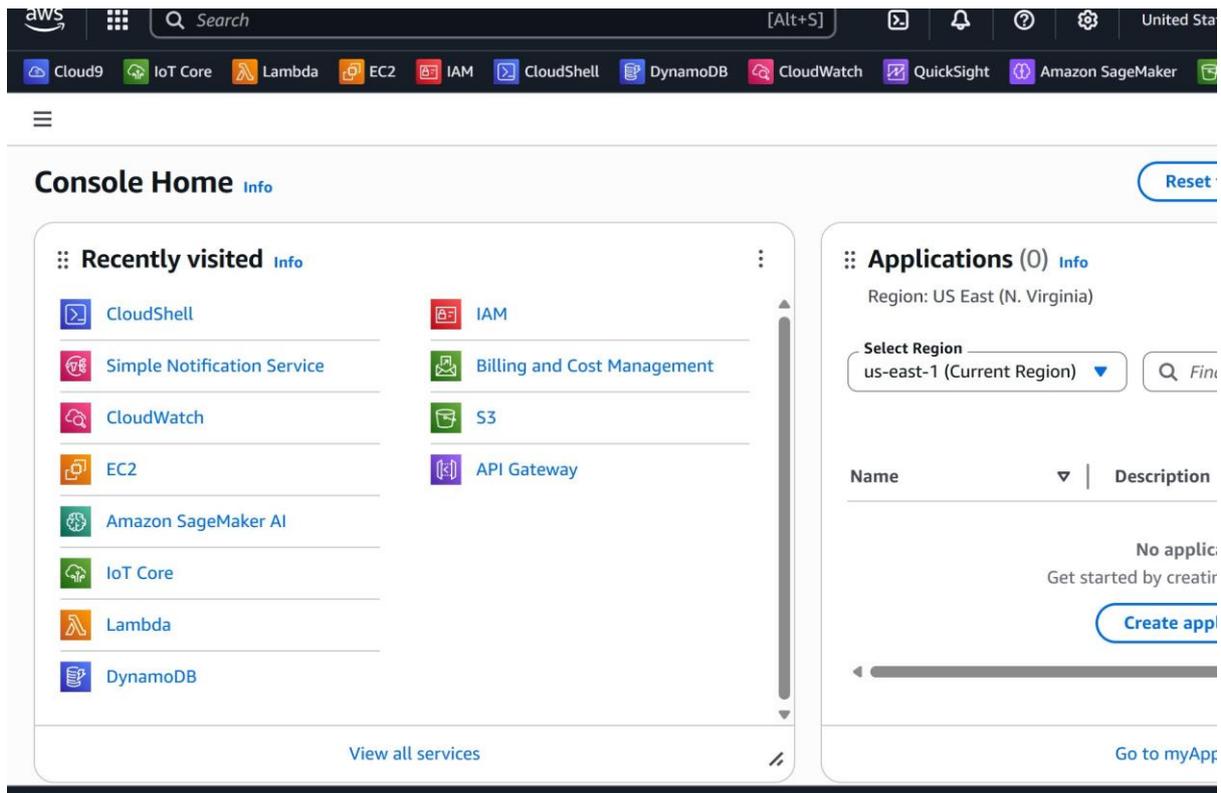


Figure 1: AWS Management Console

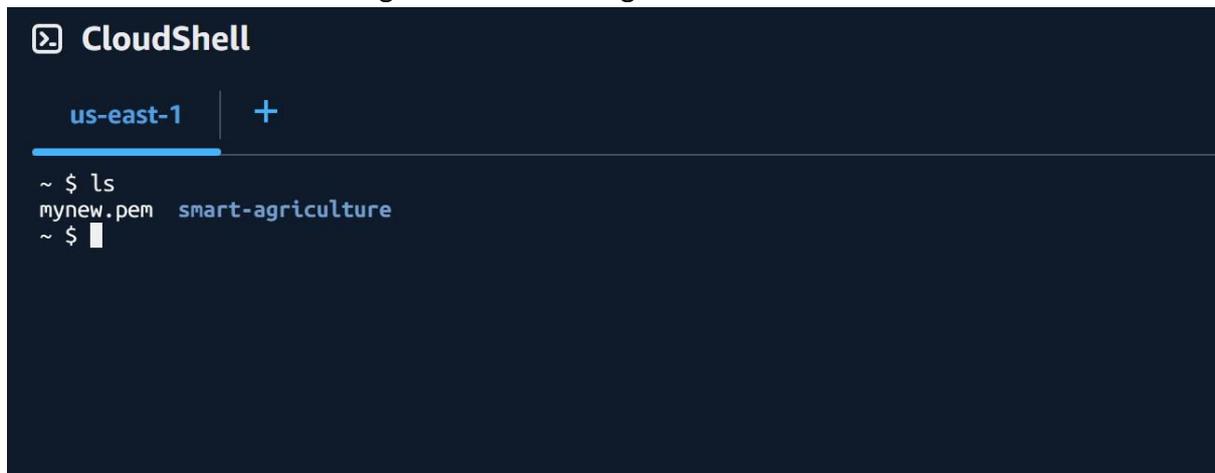


Figure 2: AWS Management Console

5 Server (EC2) Provisioning & System Setup

Amazon EC2 Installation

Sign up an EC2 instance (Ubuntu 22.04) to run the Streamlit dashboard. Opened a security group port 8501 to open up dashboard.

Set up Python packages and released the Streamlit application to monitor in realtime IoT and prediction.

Sign up an EC2 instance (Ubuntu 22.04) to run the Streamlit dashboard.

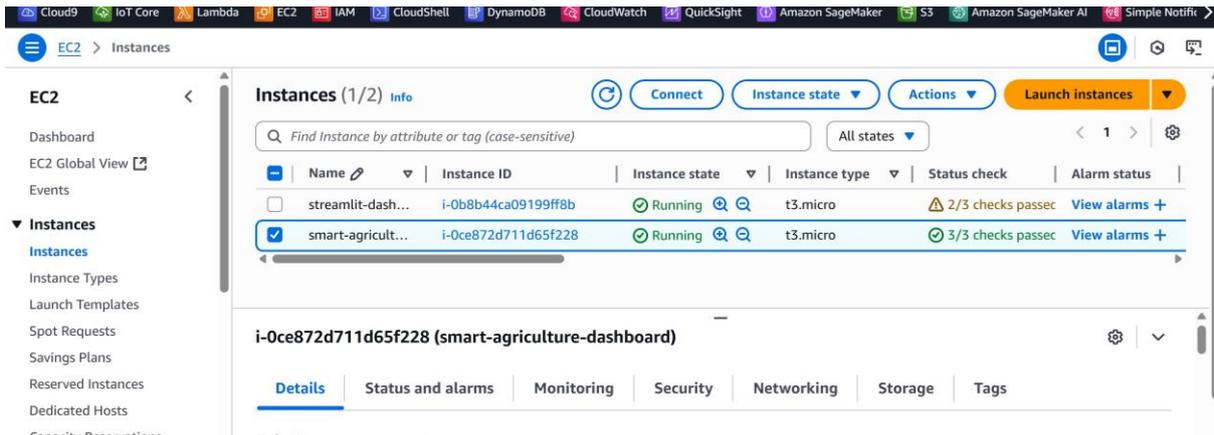


Figure 3: AWS EC2

Opened a security group port 8501 to open up dashboard.

Set up Python packages and released the Streamlit application to monitor in realtime IoT and prediction.

5.1 setup for dashboard in EC2

```
~ $ ssh -i ~/mynew.pem ec2-user@54.81.232.121
```

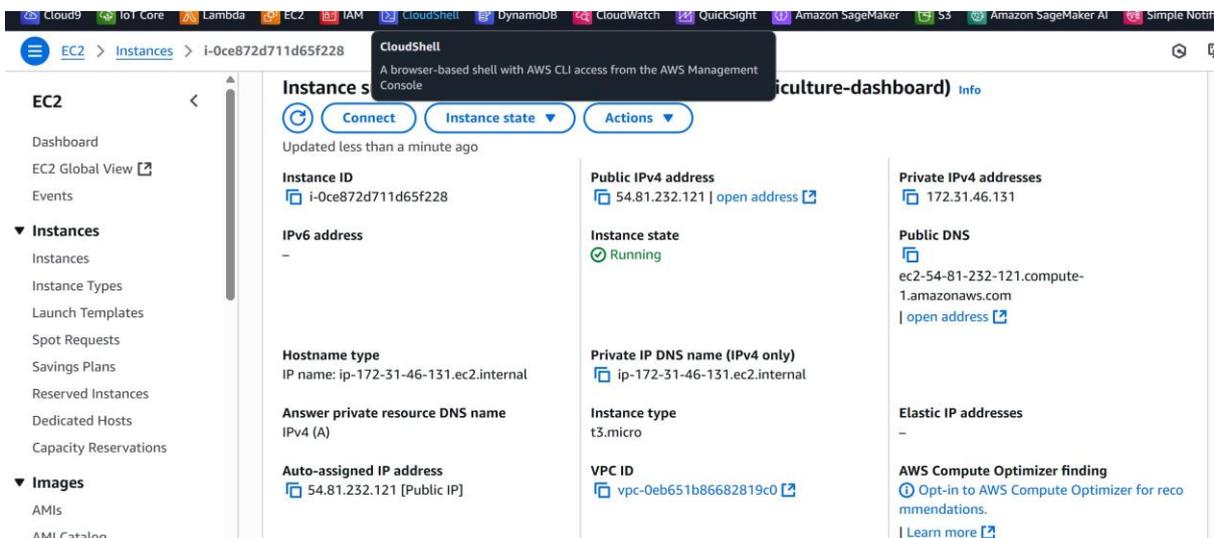


Figure 4: ec2 code

5.2 Base packages

```
[ ec2-user@ip-172-31-46-131 ~]$ ls
```

```
AmazonRootCA1.pem cert.pem dash.log dashboard.py device1_simulation nohup.out private.key
public.key simulation.log
```


Functions (2) Last fetched 0 seconds ago  **Actions**  **Create function**

Filter by attributes or search by keyword

< 1 > 

<input type="checkbox"/>	Function name 	Description 	Package type 	Runtime 	Last modified 
<input type="checkbox"/>	SaveSmartFarmData	-	Zip	Python 3.9	5 hours ago

Figure 6: AWS Management Console

SaveSmartFarmData

Function overview [Info](#) [Export to Int](#)

Diagram | **Template**

 **SaveSmartFarmData**
 Layers (0)

 **AWS IoT**

[+ Add destination](#)

[+ Add trigger](#)

Figure 7: lambda

```
import sagemaker
from sklearn . model_selection import train_test_split from sklearn . ensemble import
RandomForestClassifier from sklearn . metrics import accuracy_score , classification_report
from sagemaker . sklearn . model import SKLearnModel from sagemaker import
get_execution_role import boto3
```

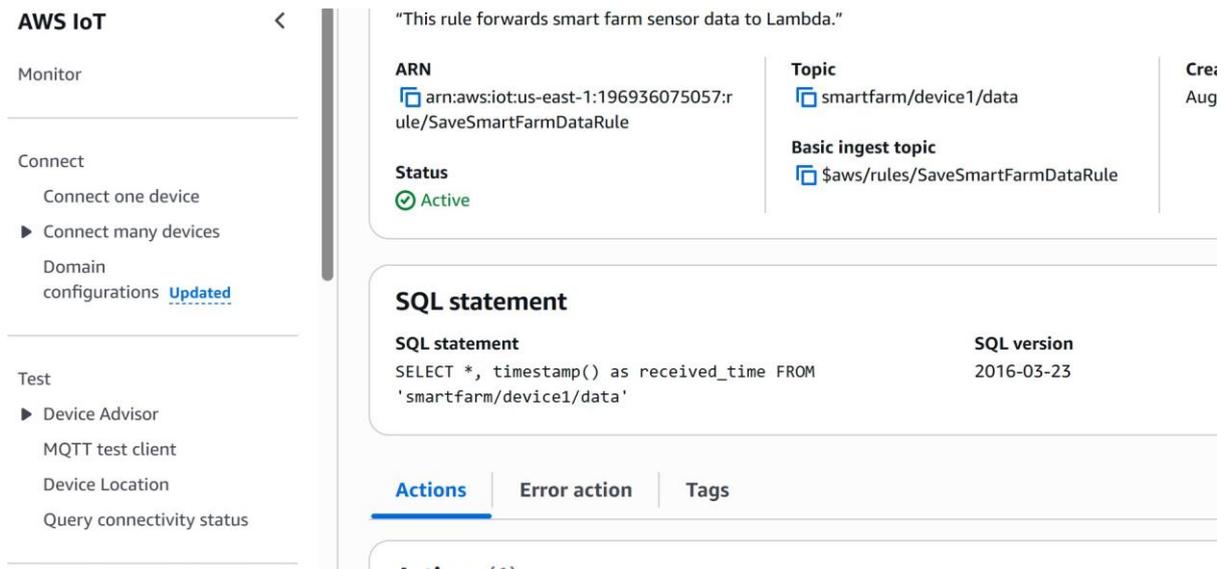


Figure 8: aws

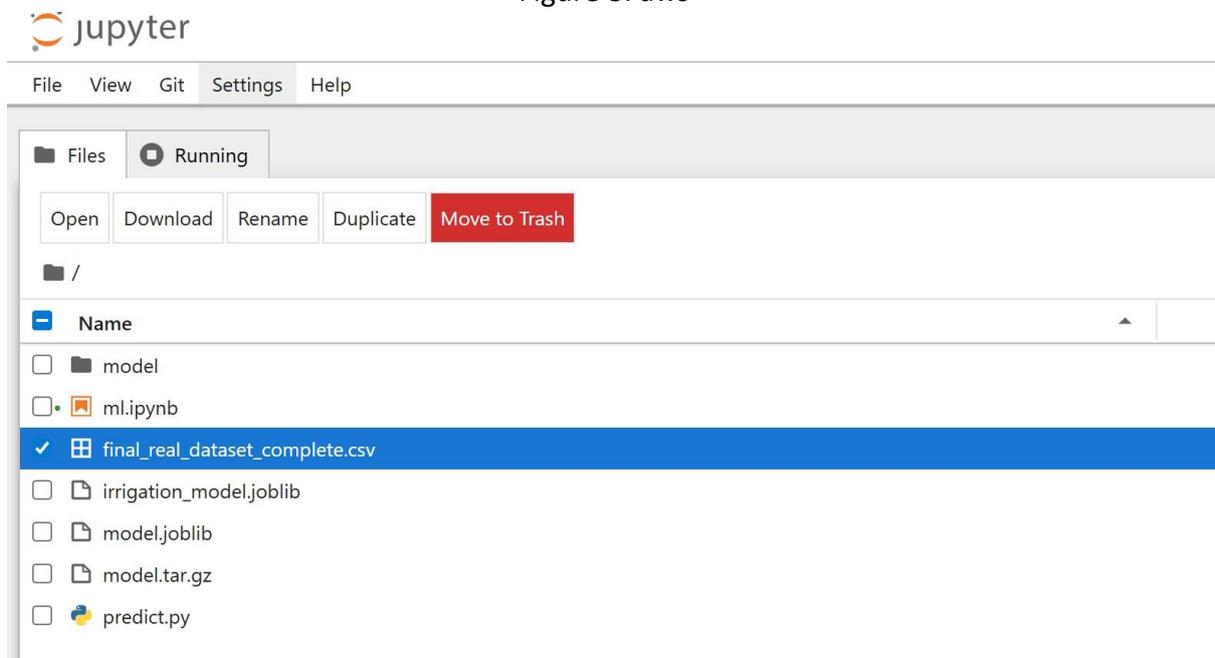


Figure 9: jupitur notebook

```
import os
```

7.1 load dataset

```
CSV_PATH = "final_real_dataset_complete . csv" # must be in working dire df = pd. read_csv(CSV_PATH)
```

```
print("Dataset_shape : " , df . shape ) print("Columns : " , df .
columns . t o l i s t ( ) )
display ( df . head ( ) )
```

```

# -----
# 2 Import Libraries
# -----
import pandas as pd
import numpy as np
import joblib
import tarfile
import sagemaker
from sklearn.model_selection import train_test_split
from sklearn.ensemble import RandomForestClassifier
from sklearn.metrics import accuracy_score, classification_report
from sagemaker.sklearn.model import SKLearnModel
from sagemaker import get_execution_role
import boto3
import os

```

Figure 10: libraries

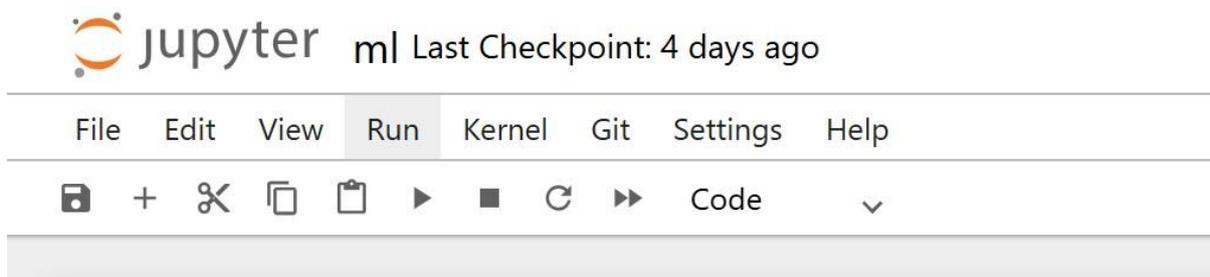


Figure 11: load

7.2 Train the Dataset

8 Deploy

```

model = RandomForestClassifier (
    n_estimators=100, random_state=42,
    class_weight='balanced'
)
model.fit(X_train, y_train)

# Evaluate y_pred = model.predict(X_test)
print("Accuracy : ", accuracy_score(y_test, y_pred)) print(" Classification _Report :\n"
, classification_report(y_test, y_pred) print("Predicted _counts :\n", pd.Series(y_pred
).value_counts())

```

	Date	Ambient_Temperature	Humidity	Rainfall	Light_Intensity	Soil_Moisture	Annual CO ₂ emissions (tonnes)
0	2024-10-03	22.240245	55.291904	0.0	556.172805	27.521109	413092654.5
1	2024-10-03	21.706763	63.949181	0.0	596.136721	14.835566	413092654.5
2	2024-10-03	21.180946	67.837956	0.0	591.124627	17.086362	413092654.5
3	2024-10-04	22.593302	58.190811	9.6	241.412476	15.336156	413092654.5
4	2024-10-04	28.929001	63.772036	9.6	444.493830	39.822216	413092654.5

Figure 12: dataset

8.1 Deploy

```
sk_model = SKLearnModel( model_data=model_data_s3 ,
    role=ROLE,
    entry_point=ENTRY_POINT_NAME,          # our robust I/O script
    framework_version=FRAMEWORK_VERSION, py_version="py3" ,
    sagemaker_session=sm_session ,
)

predictor = sk_model . deploy ( initial_instance_count=1,
    instance_type=INSTANCE_TYPE
)

endpoint_name = predictor . endpoint_name print("_Deployed_endpoint : " ,
    endpoint_name)
```

9 Final Deployment dashboard

In the cloudshell inside the EC2 bash

```
import os
import ast
import re
import time
from decimal import Decimal
```

```
import pandas as pd
import streamlit as st
```

```
# ===== Settings =====
```

```
LOG_PATH = "/home/ec2-user/simulation . log"
```

```
REFRESH_RATE = 5 # seconds between auto-refreshes
```

```
PAGE_TITLE =
    "_Smart_Farm_Dashboard"
PAGE_ICON = " "
```

```
st . set_page_config ( page_title=PAGE_TITLE, page_icon=PAGE_ICON, layout=" dashboard"
```

```
urlhttp://http://54.81.232.121:8502/ in a browser.
```

Smart Farm Dashboard

Temp (°C) 36.2 Humidity (%) 87.2 Soil Moisture (%) 18.8 Light 551

Trends

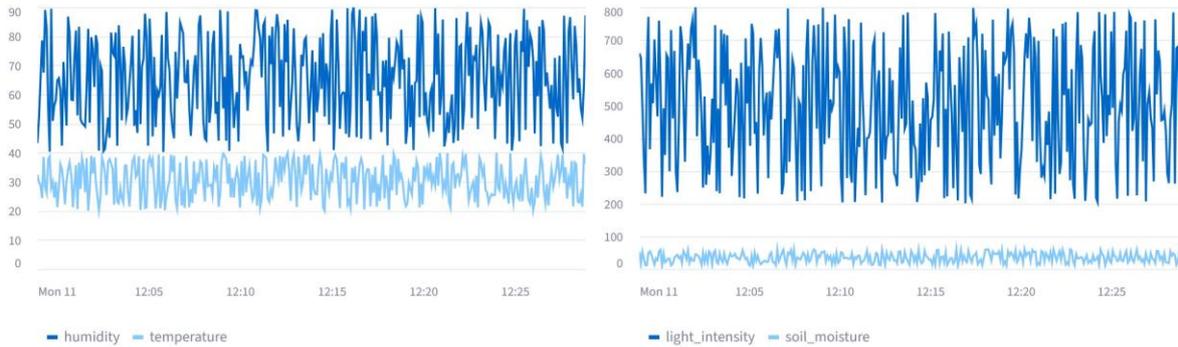


Figure 13: dashboard

```
[ ec2-user@ip -172-31-46-131 ~]$ nohup streamlit run ~/dashboard . py --server . [1] 51576
[ ec2-user@ip -172-31-46-131 ~]$ nohup: ignoring input and appending output to a file.
[1]+  Exit 1 nohup streamlit run ~/dashboard . py --server . p
[ ec2-user@ip -172-31-46-131 ~]$
```

9.1 Pushing Code to GitHub

```
# Check current branch git branch
# Stage all changes git add .
# Commit with a message git commit -m "Added
dashboard . py"
# Push to GitHub git push
origin main
```

10 References

- linux: <https://docs.kernel.org/>
- GitHub Actions: <https://docs.github.com/>
- AWS : <https://us-east-1.console.aws.amazon.com/>
- Random forest algorithm <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestRegressor.html>

- python libraries <https://docs.python.org/3/>

graphicx tikz [absolute,overlay]textpos