Demonstração Hands-On

BSC AI Factory

Classificação de Culturas Agrícolas

- 1. Criar Embeddings com Modelo Fundacional Geoespacial **Tessera**
- 2. Treinar um classificador com os embeddings
- 3. Utilizar modelo treinado para mapear/prever culturas em Portugal

BSC AI Factory

Tessera (GFM)

Sentinel-1

2 bandas x 70 imagens/ano = 140d

+

Sentinel-2

10 bandas x 70 imagens/ano = 700d

=

12 bandas x 70 imagens/ano = 840d

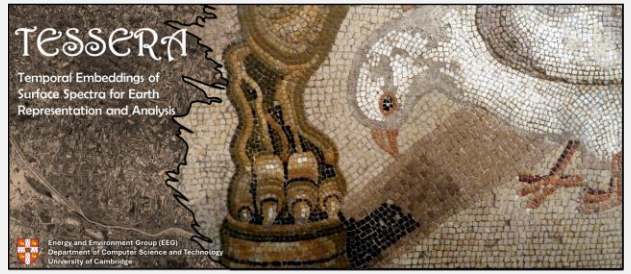

TESSERA

128d

BSC AI Factory

https://github.com/ucam-eo/tessera

Tessera (GFM)

- Encodificação a nível do pixel: foco na dimensão espectro-temporal
- Modelo Barlow-Twins: compara views temporais aleatórias de um mesmo pixel

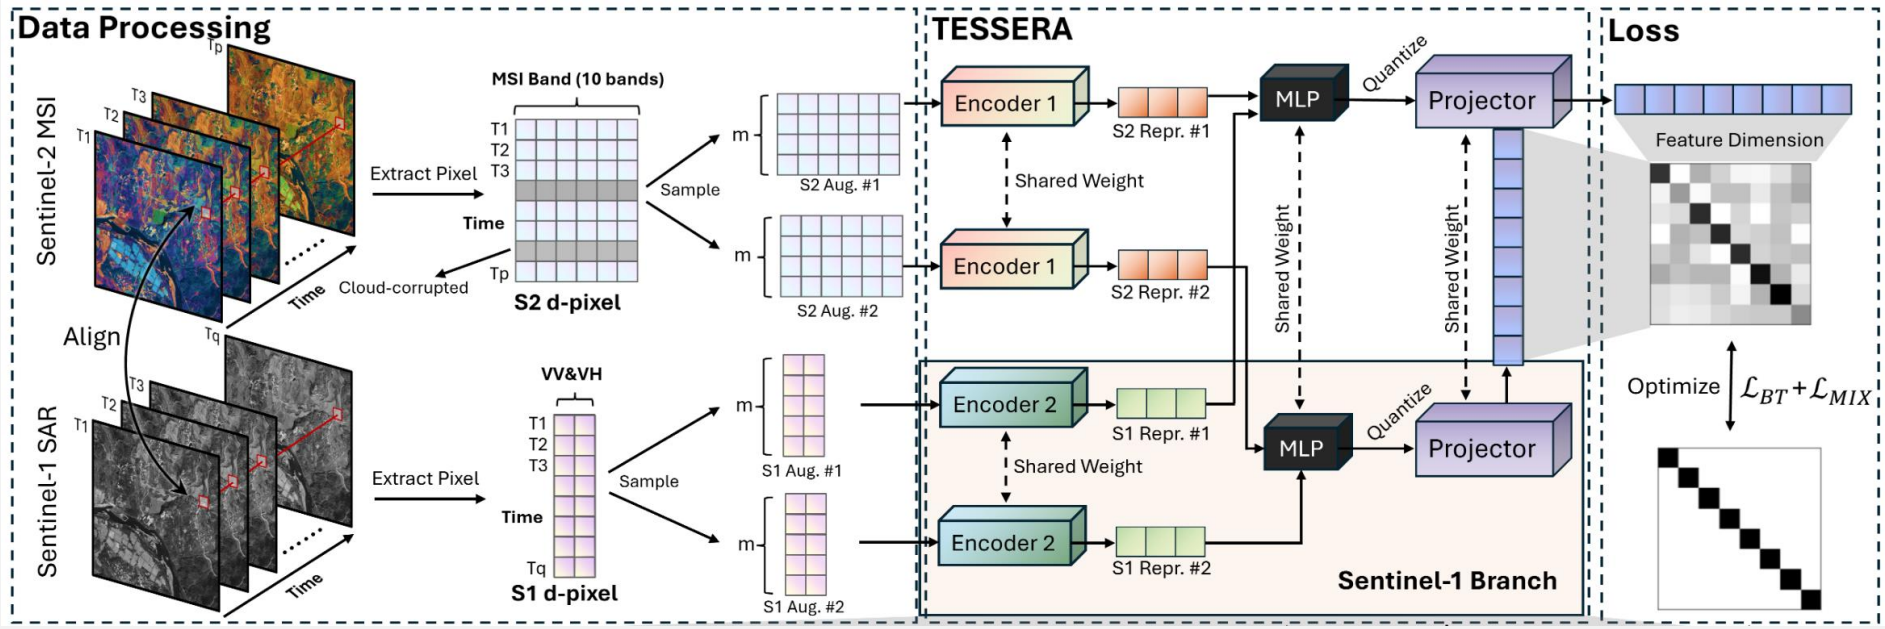

BSC AI Factory

Classificação de Culturas

- Dados de treino
  - IFAP Parcelário
  - COS
- Modelo
  - Classification Head (FCN)

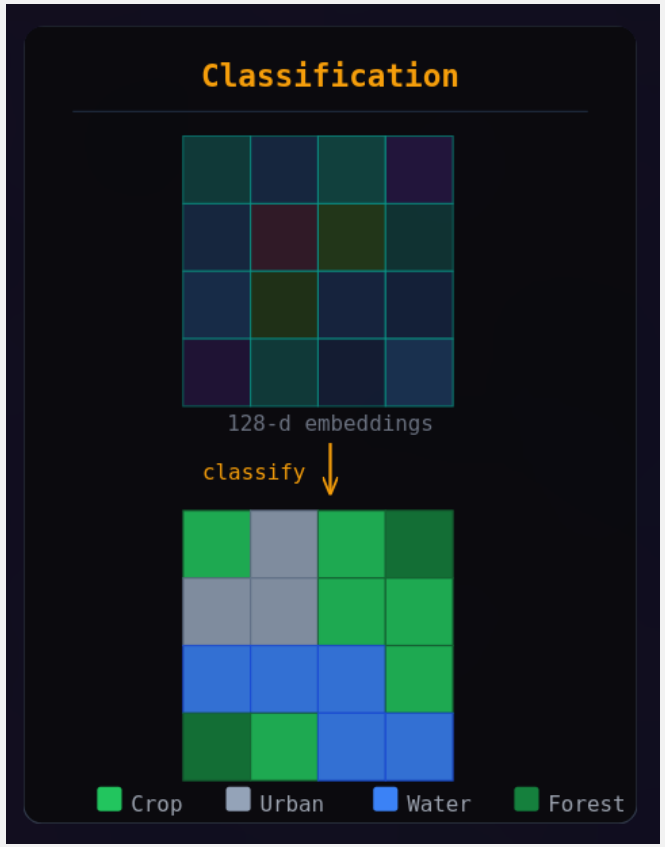

BSC AI Factory

Links Úteis

GitLab Repo

https://gitlab.a.acnca.pt/cnca/aif-pt/baif-agr-workshop

Página do Evento

https://indico.acnca.pt/event/10

BSC AI Factory

1. Abrir Shell no Open OnDemand

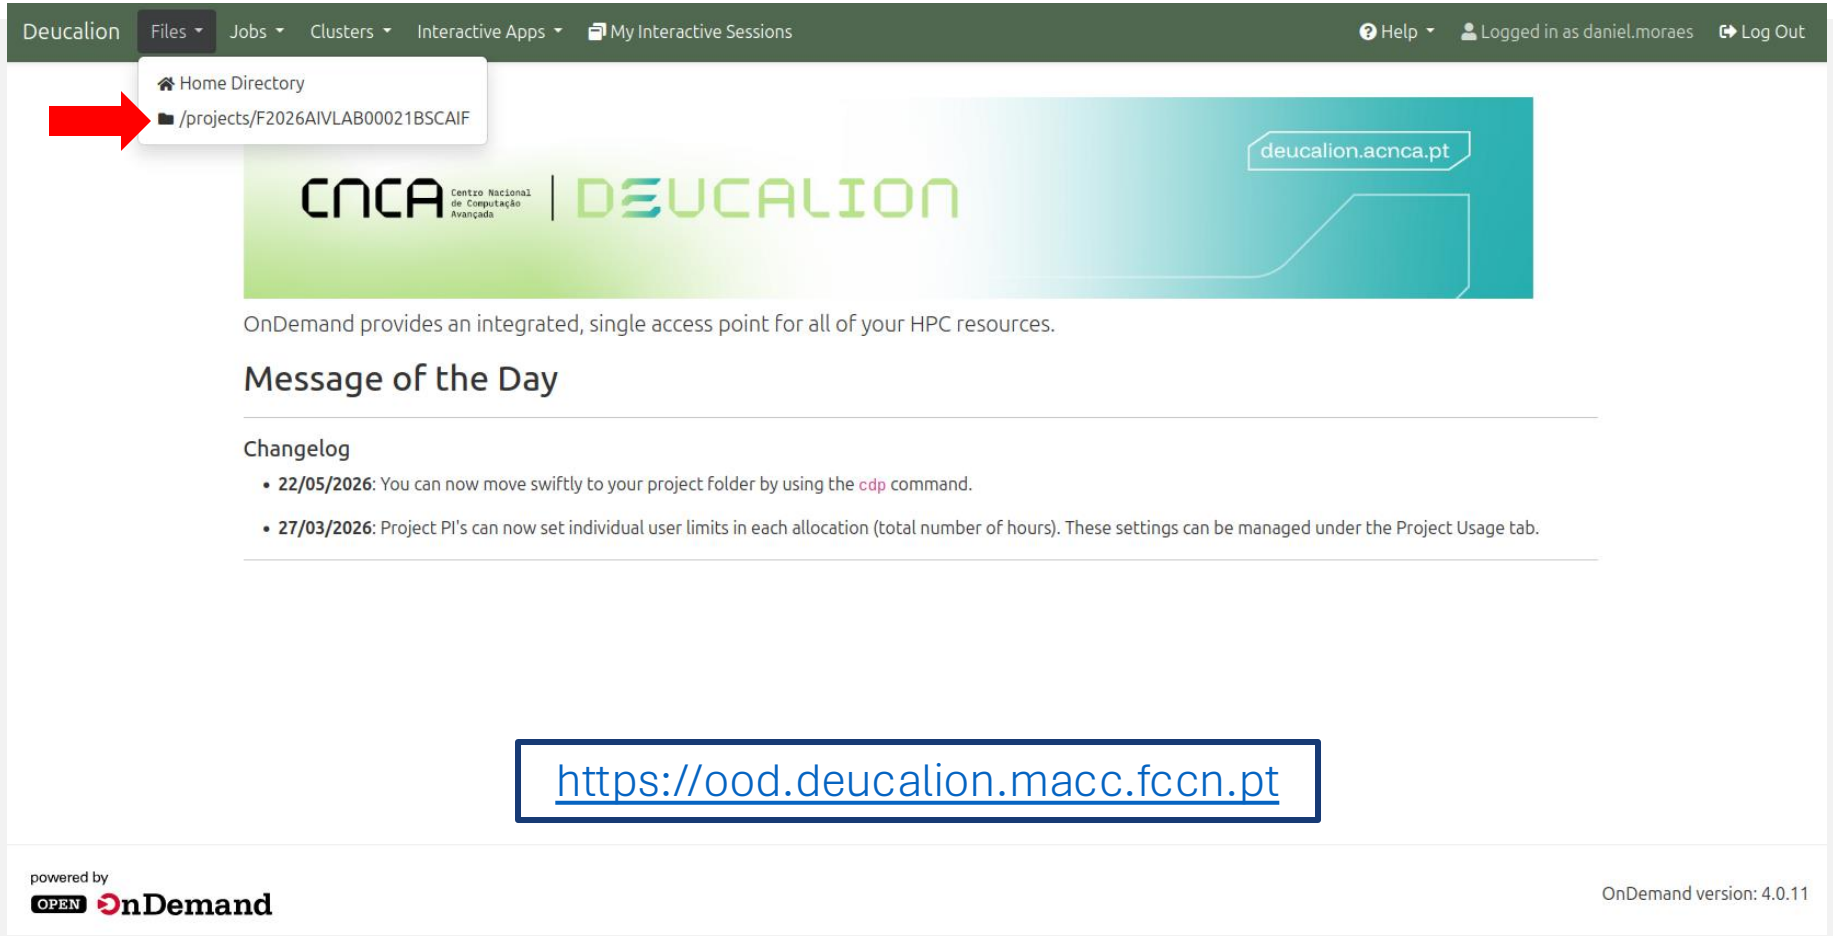

BSC AI Factory

1. Abri Shell no Open OnDemand

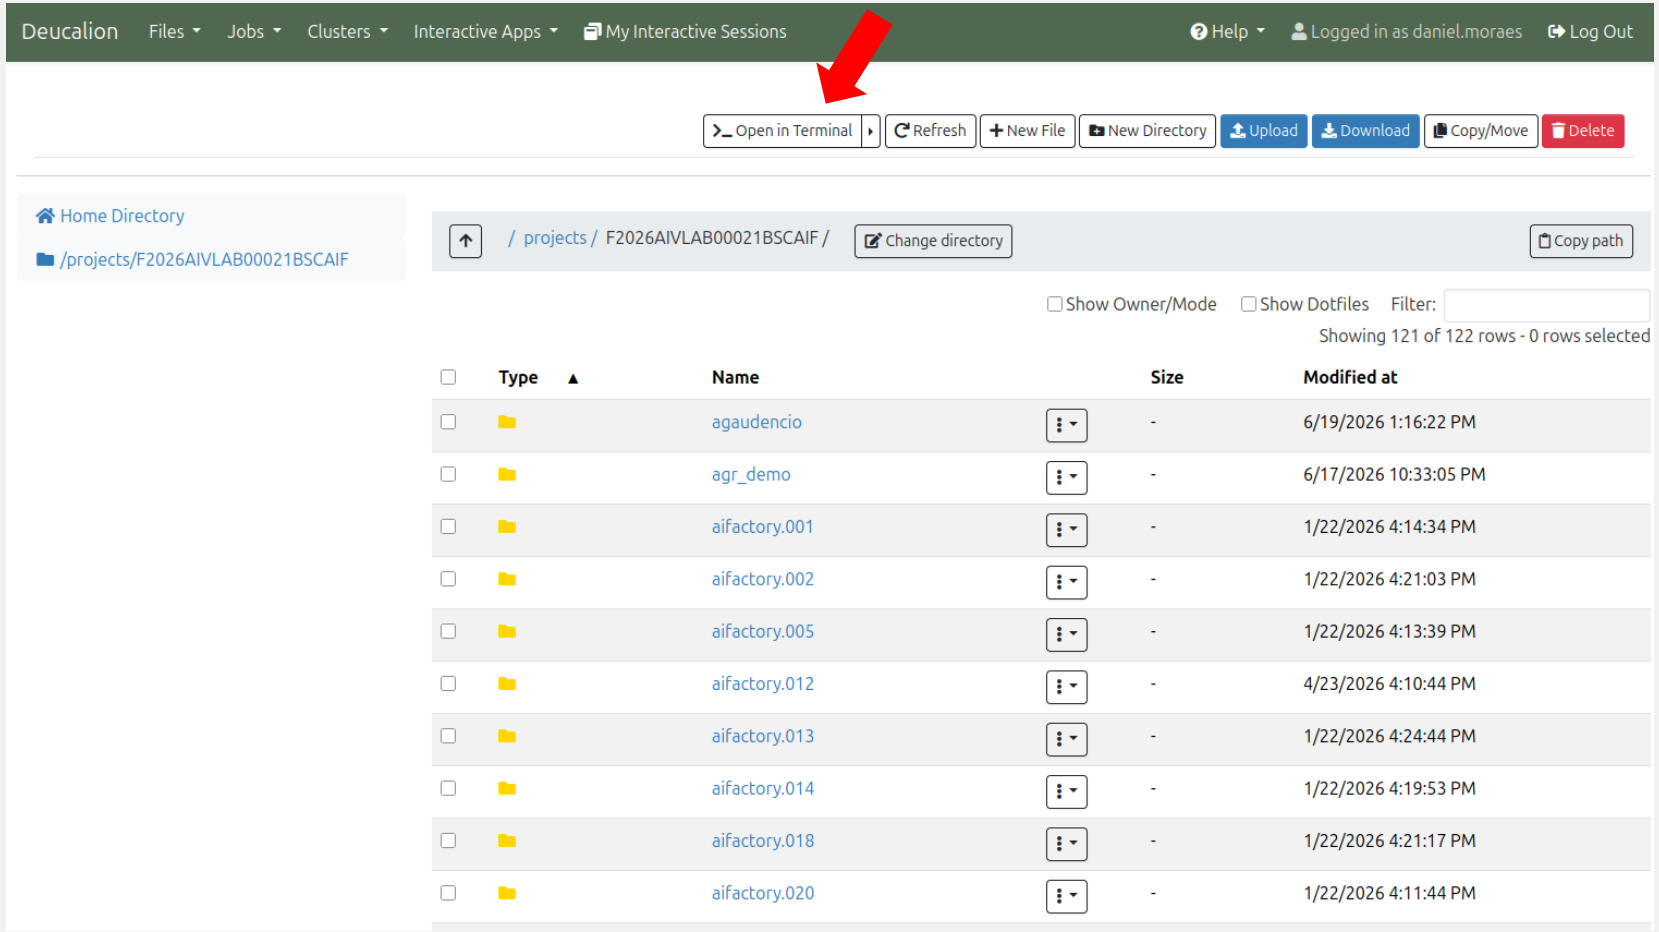

BSC AI Factory

1. Abri Shell no Open OnDemand

```
Host: login-ext      Initial directory: /projects/F2026AIVLAB00021BSCAIF

Greetings daniel.moraes

Welcome to the Portuguese EuroHPC supercomputer, *please* read documentation before use:
https://docs.macc.fccn.pt/

The following scripts are now available for user reference:

1. Home Quota Check:               quotahome
2. Lustre (projects) Quota Check:  quotaprojects
3. Slurm Billing Information:      billing

Run these scripts directly to get up-to-date information on
your storage usage and job-related billing.

[daniel.moraes@ln02 F2026AIVLAB00021BSCAIF]$
```

BSC AI Factory

2. Clone Workshop Repository

```
bash clone_agr.sh
```

```
[daniel.moraes@ln02 F2026AIVLAB00021BSCAIF]$ bash clone_agr.sh
Cloning into 'daniel.moraes/baif-agr-workshop'...
remote: Enumerating objects: 119, done.
remote: Counting objects: 100% (48/48), done.
remote: Compressing objects: 100% (47/47), done.
remote: Total 119 (delta 23), reused 0 (delta 0), pack-reused 71 (from 1)
Receiving objects: 100% (119/119), 540.01 KiB | 22.50 MiB/s, done.
Resolving deltas: 100% (24/24), done.
[daniel.moraes@ln02 F2026AIVLAB00021BSCAIF]$
```

BSC AI Factory

2. Clone Workshop Repository

- `cd $USER`
  Muda de diretório para a pasta do utilizador
- `ls`
  Lista os ficheiros e pastas existentes no diretório atual
- `cd baif-agr-workshop`
  Muda de diretório para a pasta do workshop
- `ls`

2. Clone Workshop Repository

```
[daniel.moraes@ln02 F2026AIVLAB00021BSCAIF]$ cd $USER
[daniel.moraes@ln02 daniel.moraes]$ ls
ai-assist  baif-agr-workshop  baif-llm-workshop  ml-example
[daniel.moraes@ln02 daniel.moraes]$ cd baif-agr-workshop/
[daniel.moraes@ln02 baif-agr-workshop]$ ls
assets           gpu.sh                model_checkpoints  README.md    run_train.sh     tessera  train.py
data_loader.py   interactive_map.ipynb  model.py           run_test.sh  submit_infer.sh  test.py  utils.py
[daniel.moraes@ln02 baif-agr-workshop]$
```

AI Factory

3. Abrir Jupyter Notebook

*Não feche o terminal

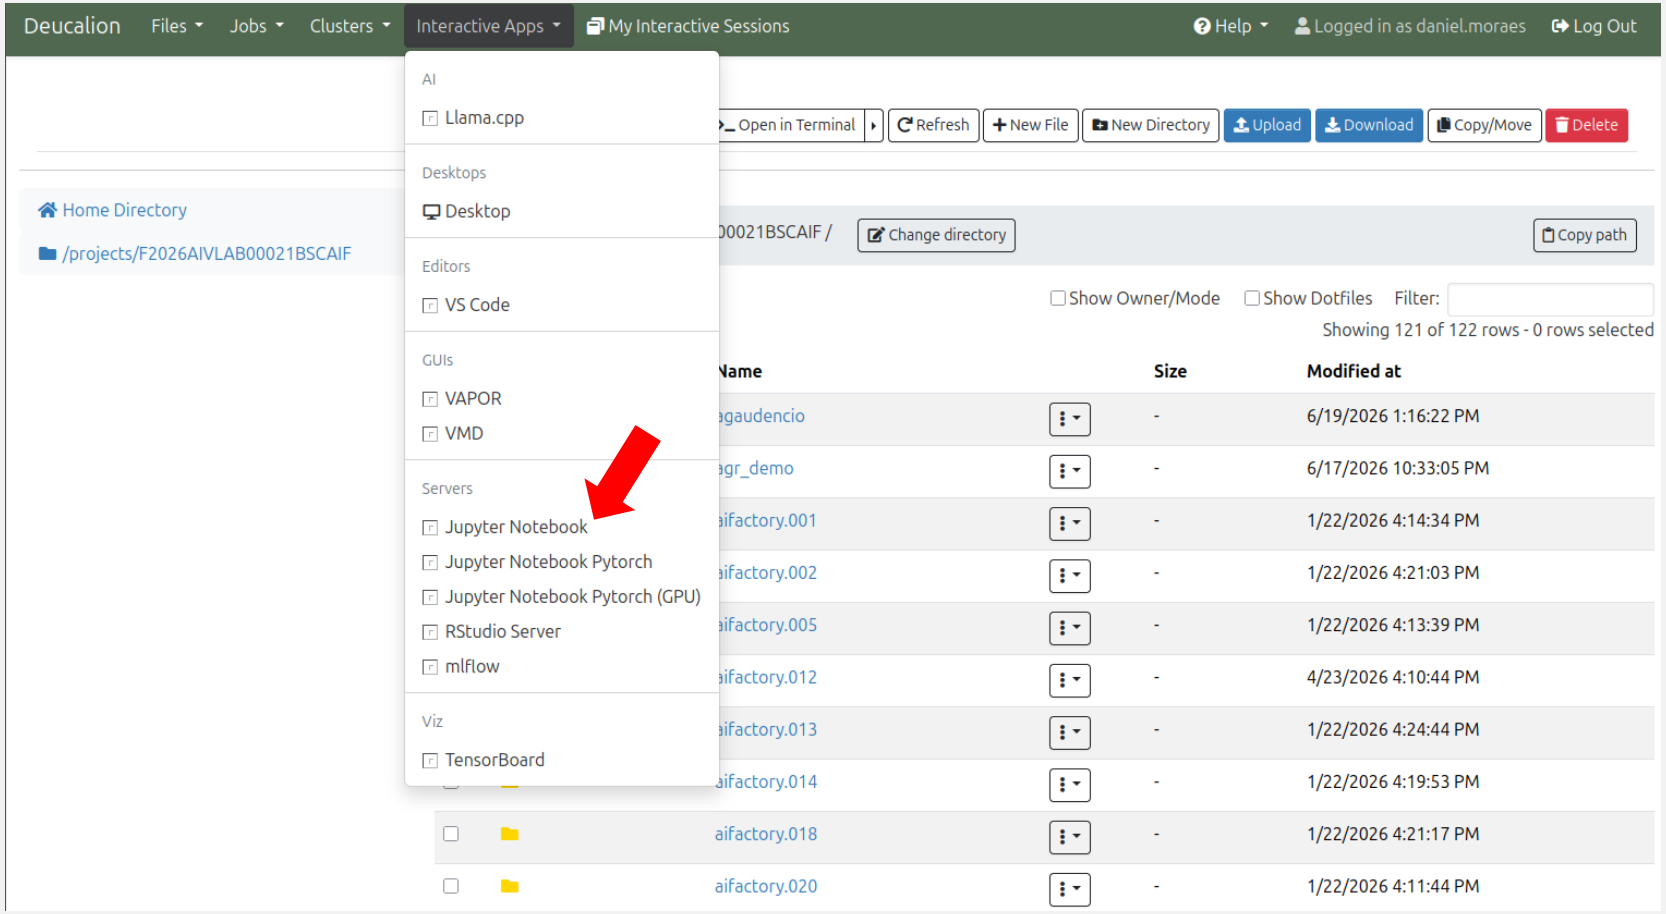

BSC AI Factory

3. Abrir Jupyter Notebook

Configurações

Nº CPU cores: 32

Nº Horas: 2

Partition: normal-x86

Account: f2026aivlab00021bscaifx

Jup. Container Image:
/projects/F2026AIVLAB00021BSCAIF/agr_demo/containers/jupyter_leafmap.sif

Reservation: aif_x86

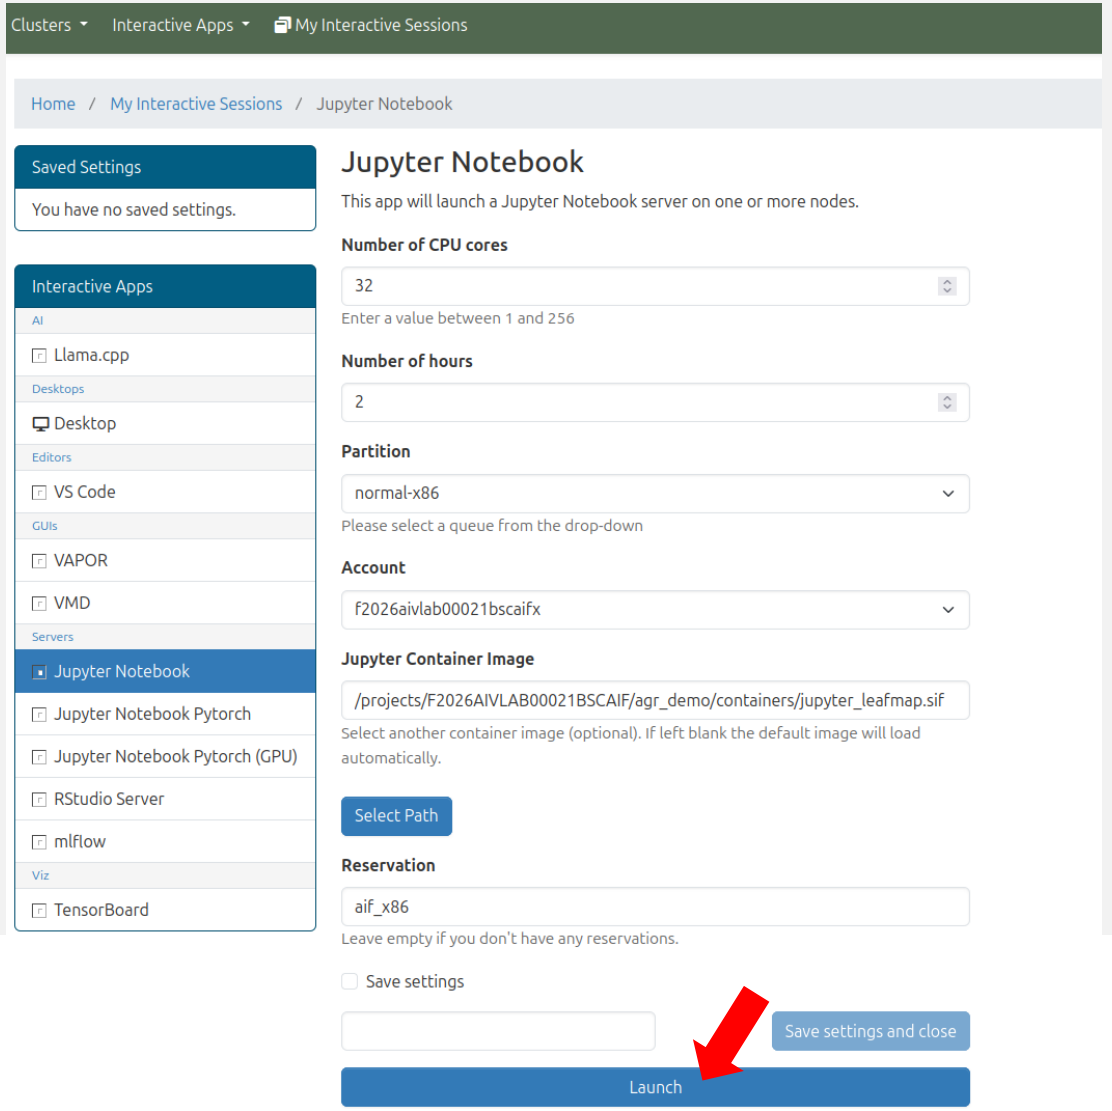

BSC AI Factory

3. Abrir Jupyter Notebook

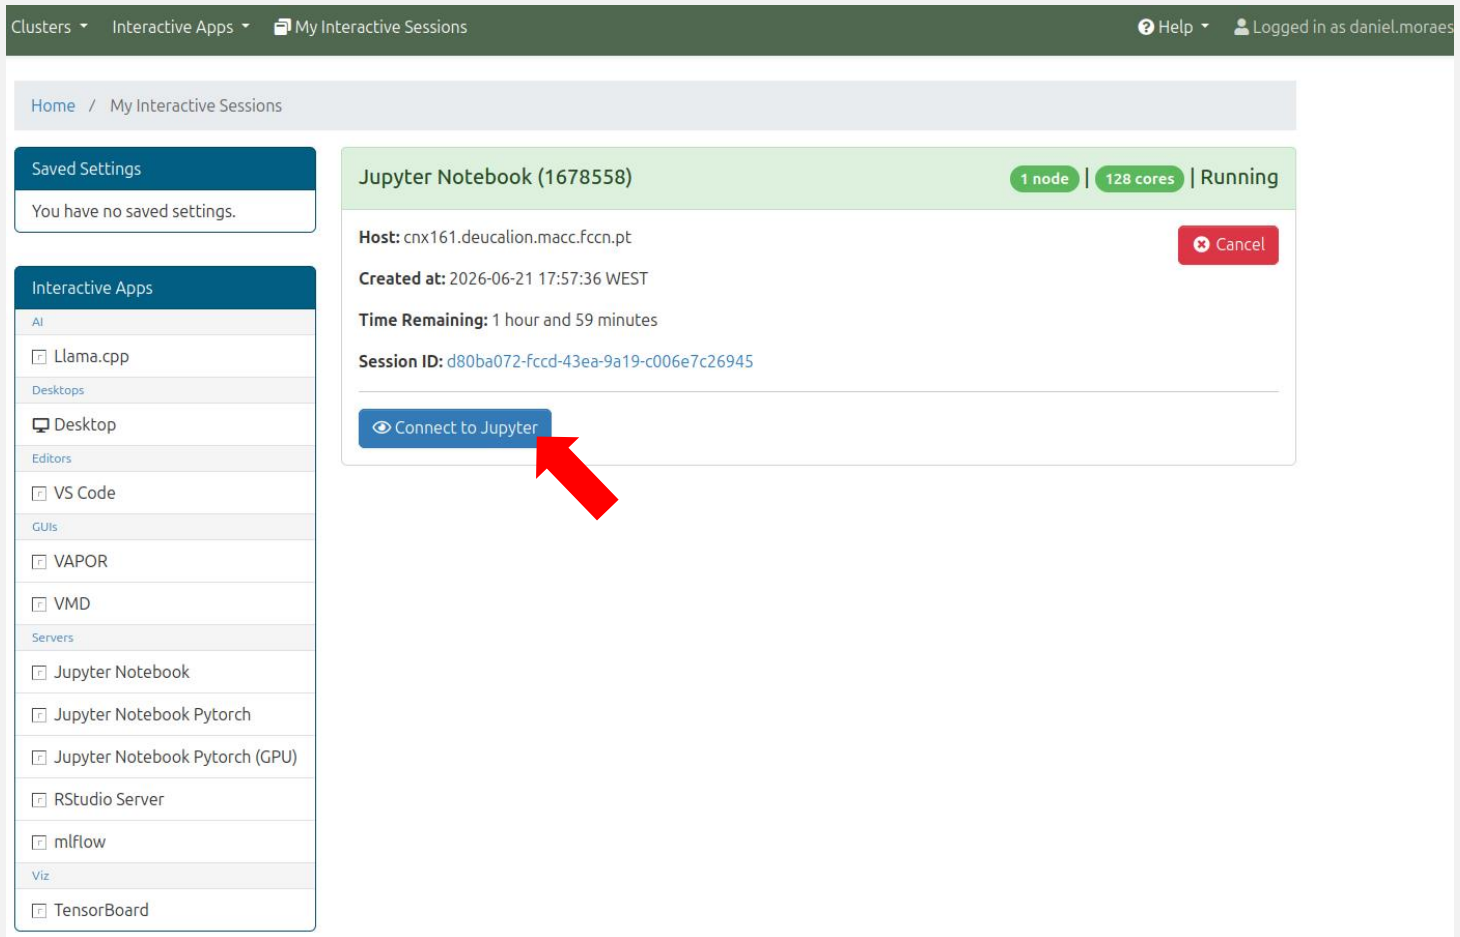

AI Factory

3. Abrir Jupyter Notebook

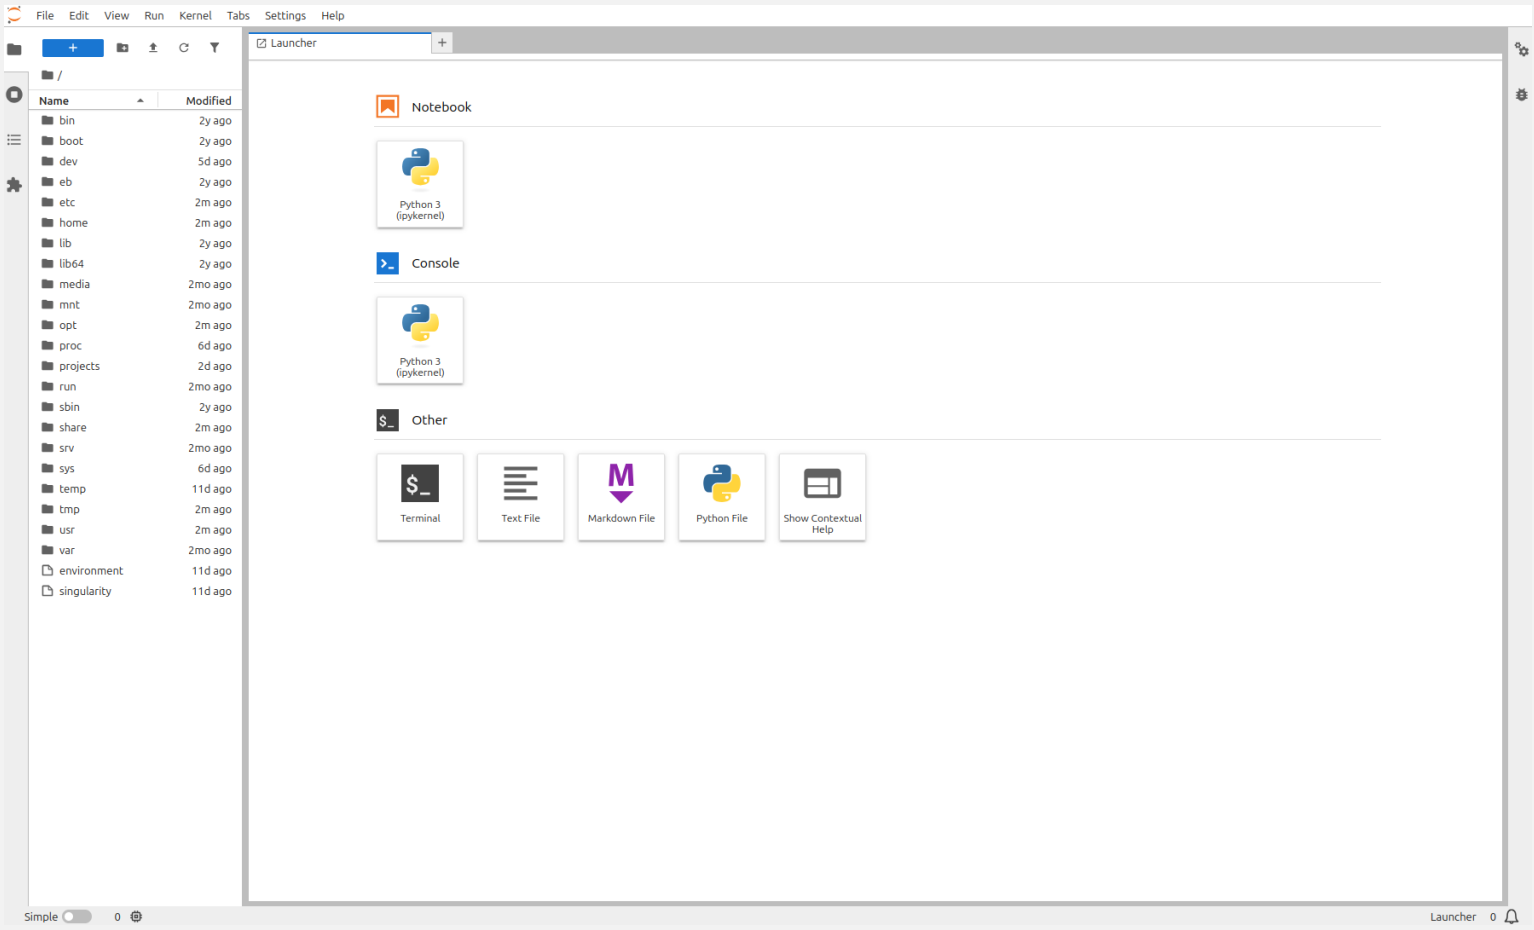

BSC AI Factory

3. Abrir Jupyter Notebook

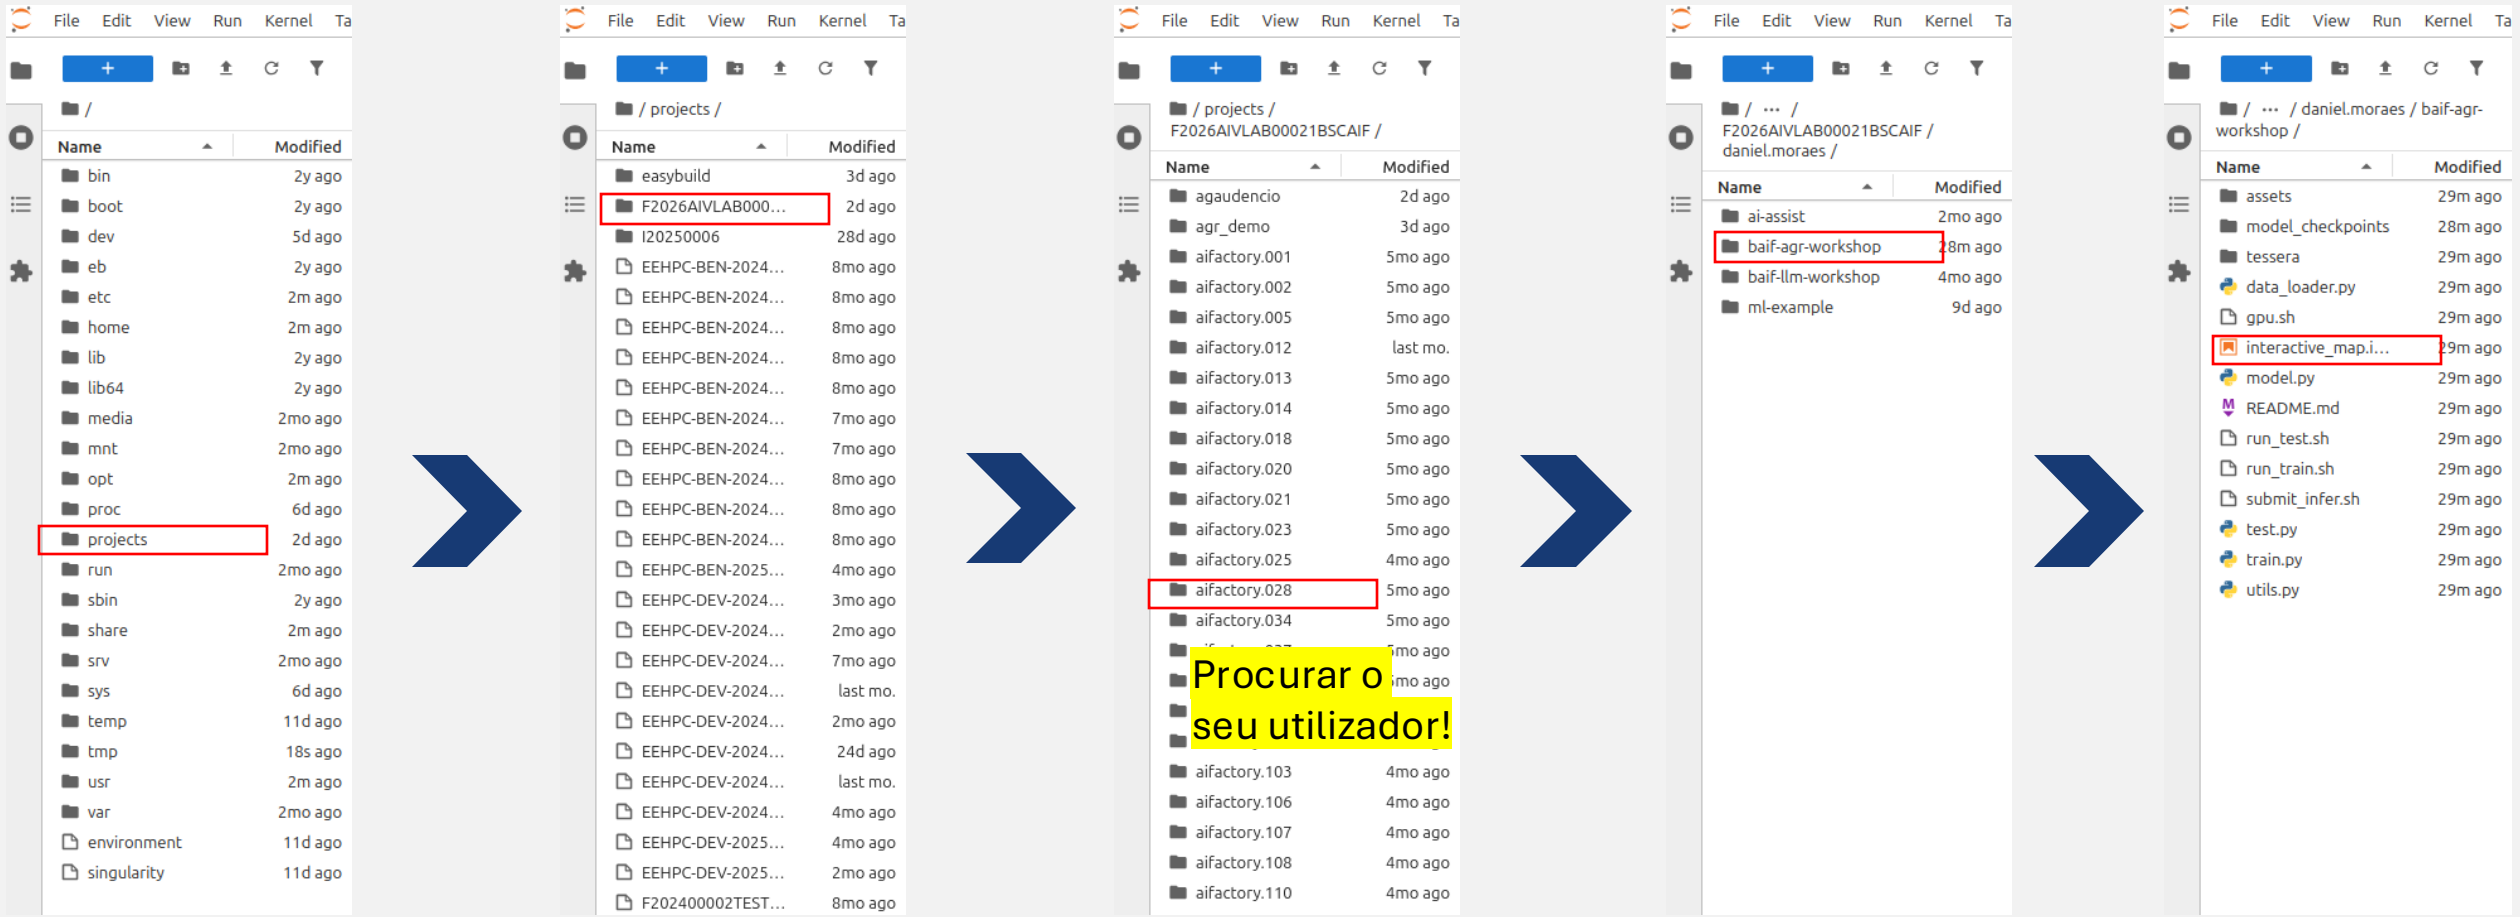

BSC AI Factory

3. Abrir Jupyter Notebook

Siga as instruções do notebook

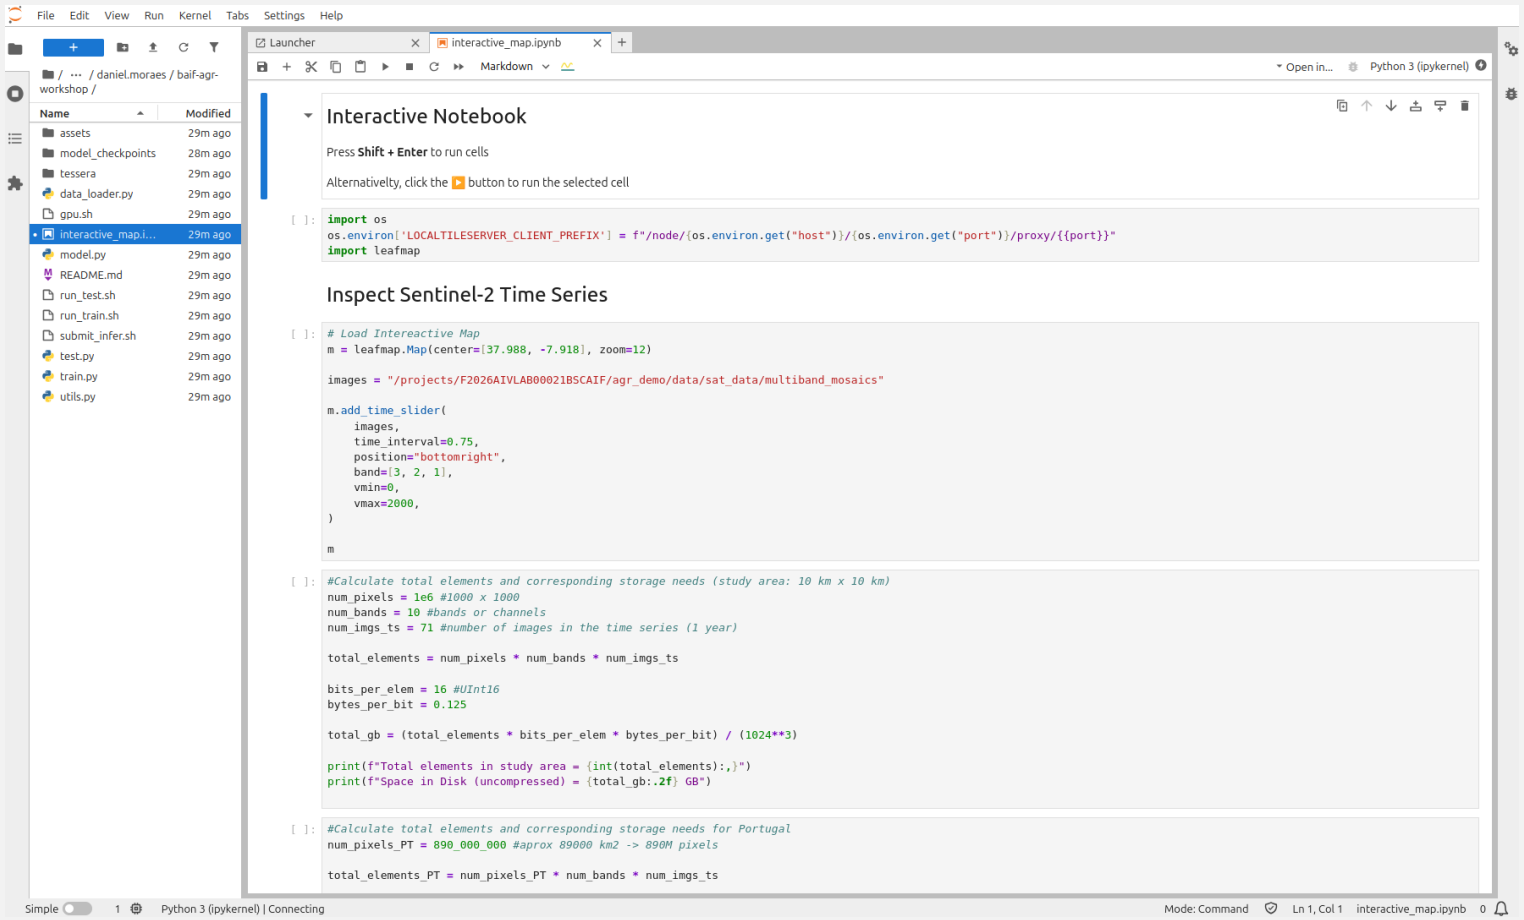

BSC AI Factory

4. Correr Job de Inferência

*De volta para o terminal. Mantenha a aba do Jupyter Notebook aberta

- `sbatch submit_infer.sh`
  Envia o job de inferência para a fila do Slurm

- `squeue --me`
  Verifica o estado do job na fila

- `tail -f slurm-infer.out`
  Verifica os outputs/logs do job
  **Faça Ctrl + C no teclado para sair do comando tail*

BSC AI Factory

4. Correr Job de Inferência

```
[daniel.moraes@ln02 baif-agr-workshop]$ sbatch submit_infer.sh
Submitted batch job 1678722
[daniel.moraes@ln02 baif-agr-workshop]$ squeue --me
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
           1678722 normal-a1    infer daniel.m  R       0:16      1 gnx522
           1678558 normal-x8 sys-dash daniel.m  R      32:01      1 cnx161
[daniel.moraes@ln02 baif-agr-workshop]$ tail -f slurm-infer.out
Running Inference...
```

Aguarde a conclusão do Job (aprox. 8 min)

BSC AI Factory

5. Visualizar Embeddings Gerados

Retornar ao Jupyter Notebook e seguir as instruções indicadas no próprio notebook.

Mantenha o terminal aberto.

BSC AI Factory

6. Correr Treino e Prediction

*De volta para o terminal. Mantenha a aba do Jupyter Notebook aberta

- `bash gpu.sh`
  Executa script que chama o comando salloc para criar sessão interativa em nó de GPU

- `bash run_train.sh`
  Executa script que dá início ao treino do modelo de classificação de culturas

Aguardar a conclusão do treino do modelo (aprox. 5 min)

BSC AI Factory

6. Correr Treino e Prediction

```
[daniel.moraes@ln02 baif-agr-workshop]$ bash gpu.sh
salloc: Pending job allocation 1678815
salloc: job 1678815 queued and waiting for resources
salloc: job 1678815 has been allocated resources
salloc: Granted job allocation 1678815
salloc: Waiting for resource configuration
salloc: Nodes gnx522 are ready for job
[daniel.moraes@gnx522 baif-agr-workshop]$ bash run_train.sh
Creating train/test datasets
Created train/test datasets
Loading model. Device: cuda
Model loaded
Starting training
Epoch [1/15], Average Loss: 0.8140, Average Loss Test: 0.5781, Time/epoch: 0.36min
Epoch [2/15], Average Loss: 0.6390, Average Loss Test: 0.5445, Time/epoch: 0.3min
```

6. Correr Treino e Prediction

*De volta para o terminal. Mantenha a aba do Jupyter Notebook aberta

- `bash run_test.sh`
  Executa script que usa o modelo treinado para prever as culturas em PT

- Aguardar a conclusão da prediction (aprox. 5 min)

- `exit`
  Encerra a sessão interativa no nó de GPU

BSC AI Factory

6. Correr Treino e Prediction

```
[daniel.moraes@gnx522 baif-agr-workshop]$ bash run_test.sh
Loading model. Device: cuda
Model loaded
Running predictions: Batch 0/143
Running predictions: Batch 1/143
Running predictions: Batch 2/143
Running predictions: Batch 3/143
Running predictions: Batch 4/143
Running predictions: Batch 5/143
Running predictions: Batch 6/143
Running predictions: Batch 7/143
Running predictions: Batch 8/143
Running predictions: Batch 9/143
Running predictions: Batch 10/143
```

7. Visualizar Mapa de Culturas

Retornar ao Jupyter Notebook e seguir as instruções indicadas no próprio notebook.

BSC AI Factory

Fim
Obrigado

daniel.moraes@acnca.pt

Co-funded by the European Union

EuroHPC Joint Undertaking

Generalitat de Catalunya

Plan de Recuperación, Transformación y Resiliencia

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101234399 and Spain, Portugal, Romania and Türkiye. The JU receives support from the European Union's Horizon Europe Programme.

BSC AI Factory Partners

Affiliated entities

Centro Nacional de Computação Avançada