

Machine learning and deep learning services for the EOSC

The DEEP-Hybrid-DataCloud project

EGI Conference 2020

November 2 – 5, 2020

Álvaro López García
aloga@ifca.unican.es
Spanish National Research Council

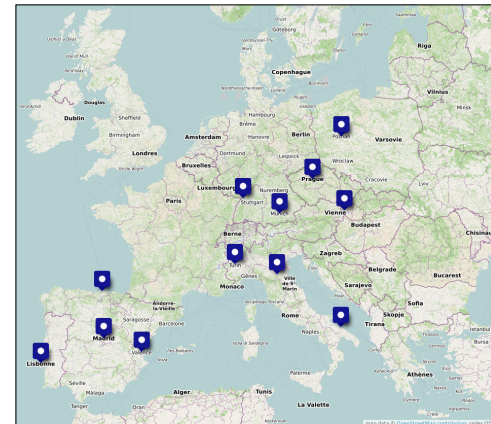


DEEP-HybridDataCloud has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 777435.



DEEP project in 1 slide

- Designing and **Enabling E-Infrastructures** for intensive data **Processing** in a **Hybrid DataCloud** (Grant agreement number 777435, Nov 2017 – Apr 2020)
- **Global objective:** Promote the use of **intensive computing services** by different research communities and areas, and the **support by the corresponding e-Infrastructure** providers and open source projects
 - Focusing on **Machine learning, Deep learning, and Post processing** services for the European Open Science Cloud



Context: the EOSC

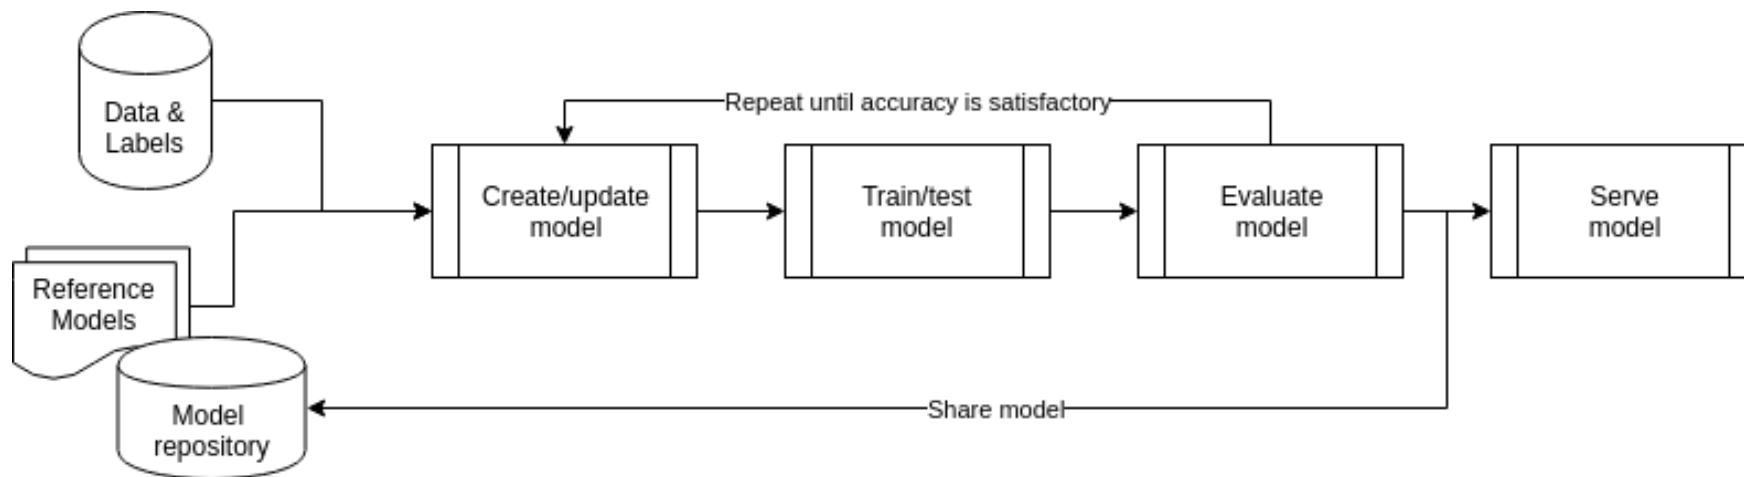
- The **European Open Science Cloud** was launched in 2016 as part of the European Cloud Initiative
- The EOSC is a **digital platform to support and develop science** and open innovation
- Europe's virtual environment for all researchers
- Providing seamless access to data and interoperable services across borders
- Access is delivered through the **EOSC portal**
- Launched in 2018, continuously being enhanced and developed



- Provide **specialized cloud services** to develop, exploit and share machine learning and deep learning applications → **Service-Oriented Architectures and platforms**
 - Covering the whole machine learning application development cycle
 - Focused on all types of users and user knowledge
- Transparent **access to specialized computing resources** (accelerators, high performance computing) → **reduce entry barrier**
- Build an **EOSC machine learning marketplace** as an application exchange → **ease of use, foster collaboration, knowledge dissemination**
- **Run the same application everywhere** (latptop, cloud, HPC, etc.) → **ease of use**



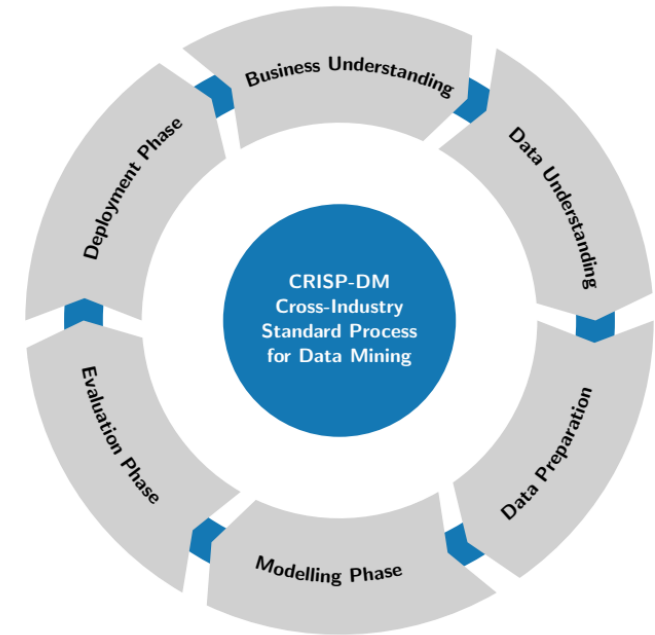
Machine learning development cycle



- The project covers **all development phases** of a machine learning application:
 - Creation, update, improvement of a model (existing or new)
 - Model training, test, evaluation
 - Model deployment as a service
 - Model publication for sharing and reuse

The CRISP-DM cycle

- We are focusing on several phases of the Cross-Industry Standard Process for Data Mining (CRISP-DM) cycle:
 - **Modeling, evaluation and deployment** phases
- Provide a **training and evaluation platform**
- Provide a **model deployment service**
- Provide a **catalog to share applications and models**



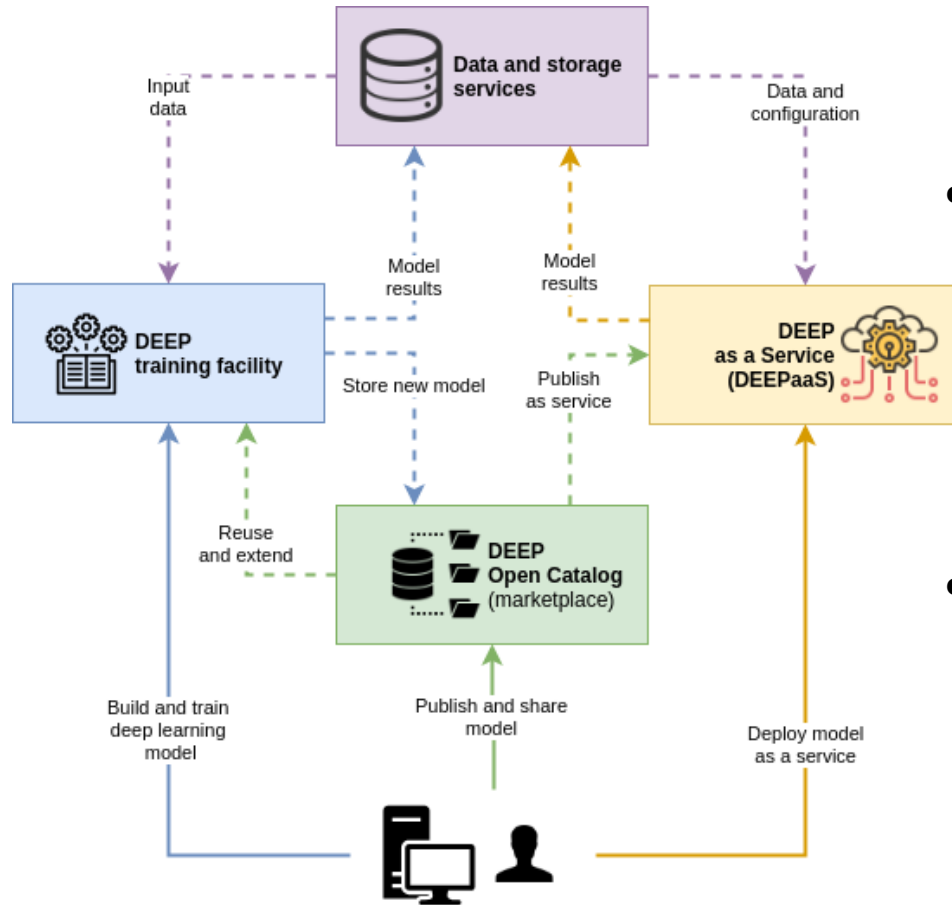
- **Category 1:** Deploy a readily trained network for somebody else to use it on his/her data set
 - Domain knowledge
- **Category 2:** Retrain (parts of) a trained network to make use of its inherent knowledge and to solve a new learning task
 - Domain + machine learning knowledge
- **Category 3:** Completely work through the deep learning cycle with data selection, model architecture, training and testing
 - Domain + machine + technological knowledge

Docker containers

- DEEP leverages Docker containers to encapsulate the user application, together with the required environment
- Containers are automatically built and published to the marketplace
- Docker allows users to run the same application everywhere:
 - Using Docker, if available
 - Using uDocker (user space tool, no system-wide install needed) if Docker is not available (like in HPC environments)
 - See: <https://github.com/indigo-dc/udocker/>
 - Using serverless components (like Apache OpenWhisk)



The DEEP services



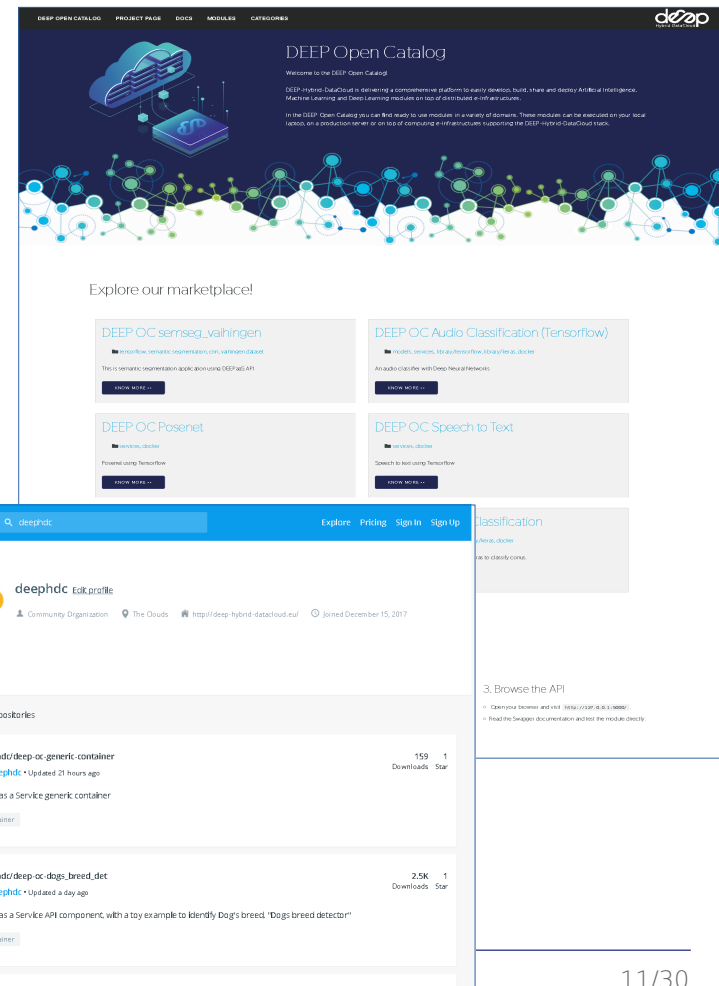
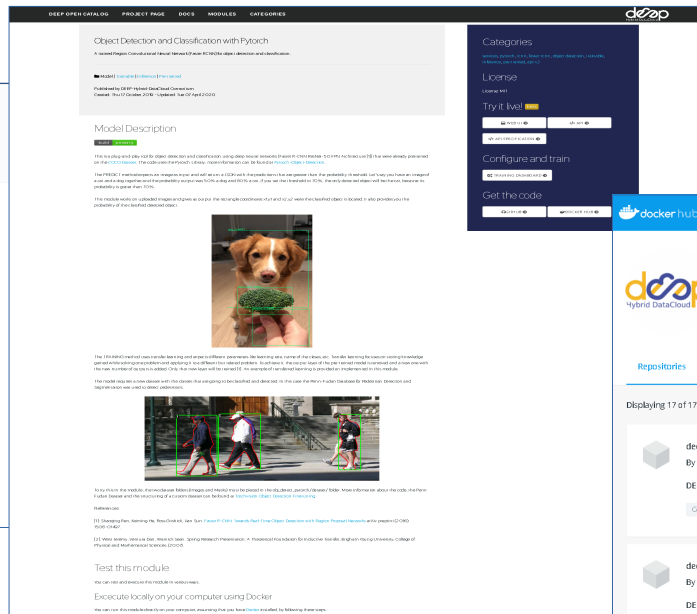
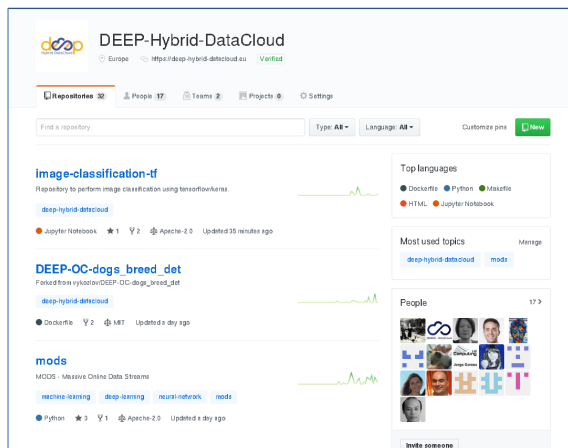
- High level specialized services for exploitation through EOSC
 - **DEEP Open Catalog** → sharing
 - **DEEP training facility** → training
 - **DEEP as a Service** → deployment
- Integration with storage from external initiatives (e.g. eXtreme-DataCloud, EGI Data Hub, etc.)

DEEP Open Catalog



Sharing knowledge, models and applications between users

<https://marketplace.deep-hybrid-datacloud.eu>



Browsing the marketplace

DEEP OPEN CATALOGPROJECT PAGEDOCSMODULESCATEGORIES

deepHybrid DataCloud

Explore our marketplace!

Artistic style transfer

Model | Trainable | Inference | Pre-trained

A module to apply artistic style transfer using pytorch.

KNOW MORE >>

Bird sound classifier

Model | Trainable | Inference | Pre-trained

Classify audio files among bird species from the Xenocanto dataset.

KNOW MORE >>

TF Benchmarks

Model | Trainable

tf_cnn_benchmarks accessed via DEEPaaS API

KNOW MORE >>

Object Detection and Classification with Pytorch

Model | Trainable | Inference | Pre-trained

A trained Region Convolutional Neural Network (Faster RCNN) for object detection and classification.

KNOW MORE >>

2D semantic segmentation

Model | Trainable | Inference | Pre-trained

2D semantic segmentation trained on the Vaihingen dataset

KNOW MORE >>

Train an audio classifier

Model | Trainable | Inference | Pre-trained

Train your own audio classifier with your custom dataset. It comes also pretrained on the 527 AudioSet classes.

KNOW MORE >>

Marketplace entry

DEEP OPEN CATALOG

PROJECT PAGE

DOCS

MODULES

CATEGORIES

deep
Hybrid DataCloud

Object Detection and Classification with Pytorch

A trained Region Convolutional Neural Network (Faster RCNN) for object detection and classification.

Model | [Trainable](#) | [Inference](#) | [Pre-trained](#)

Published by DEEP-Hybrid-DataCloud Consortium
Created: Thu 17 October 2019 - Updated: Mon 15 June 2020

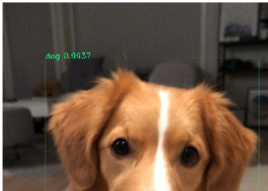
Model Description

build passing

This is a plug-and-play tool for object detection and classification using deep neural networks (Faster R-CNN ResNet-50 FPN Architecture [1]) that were already pretrained on the [COCO Dataset](#). The code uses the Pytorch Library, more information can be found at [Pytorch-Object-Detection](#).

The PREDICT method expects an image as input and will return a JSON with the predictions that are greater than the probability threshold. Let's say you have an image of a cat and a dog together and the probability output was 50% a dog and 80% a cat, if you set the threshold to 70%, the only detected object will be the cat, because its probability is grater than 70%.

This module works on uploaded images and gives as output the rectangle coordinates x1,y1 and x2,y2 were the classified object is located. It also provides you the probability of the classified detected object.



Categories

services, pytorch, rcnn, faster-rcnn, object detection, trainable, inference, pre-trained, api-v2

License

License: MIT

Try it live!

WEB UI

API

API SPECIFICATION

Configure and train

TRAINING DASHBOARD

Get the code

GITHUB

DOCKER HUB

Local inference and prediction

```
(udocker) alvaro:~ $ udocker pull deephdc/deep-oc-obj_detect_pytorch
Info: downloading layer sha256:2746a4a261c9e18bfd7ff0429c18fd7522acc14fa4c7ec8ab37ba5ebaadbc584
(...)
(udocker) alvaro:~ $ udocker create deephdc/deep-oc-obj_detect_pytorch
(udocker) alvaro:~ $ udocker run -p 5000:5000 deephdc/deep-oc-obj_detect_pytorch
Warning: this container exposes TCP/IP ports
Warning: non-existing user will be created
```

```
*****
*
*          STARTING 9bb8179d-c731-348c-bbf4-e43045ee8936
*
*****
executing: deepaas-run
```

```
##      ###
##      #####
.#####. .#####.
##  ## //  ## //  ##  ##
##. .##  ## //  ##  ##
##  ##  ####  ####  #####
      Hybrid-DataCloud  ##
```

Welcome to the DEEPaaS API API endpoint. You can directly browse to the API documentation endpoint to check the API using the builtin Swagger UI or you can use any of our endpoints.

```
API documentation: http://0.0.0.0:5000/ui
API specification: http://0.0.0.0:5000/swagger.json
V2 endpoint: http://0.0.0.0:5000/v2
```


Local inference and prediction

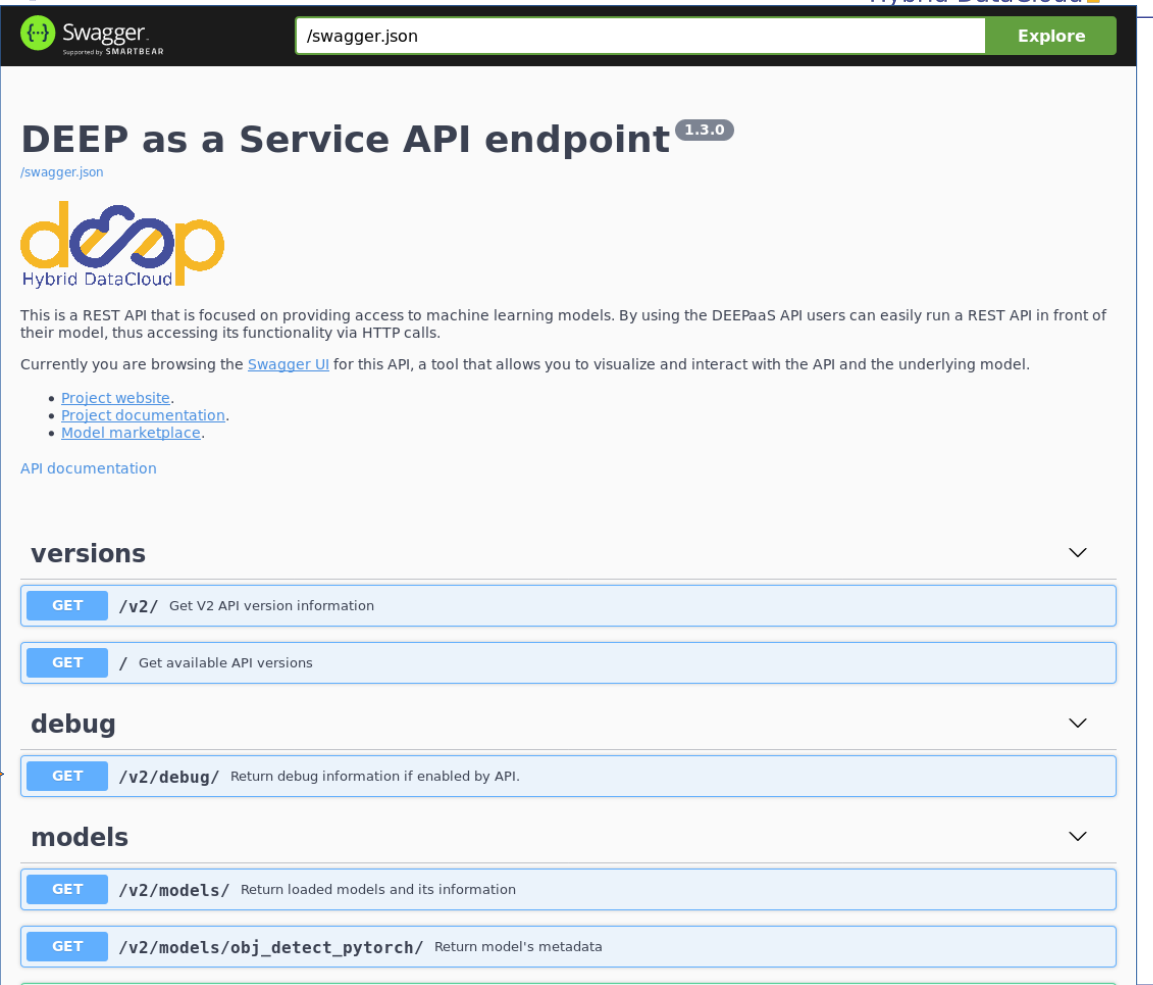
```
(udocker) alvaro:~ $ udocker pull deephdc/deep-oc-obj_detect
Info: downloading layer sha256:2746a4a261c9e18bfd7ff0429c18f
(...)
(udocker) alvaro:~ $ udocker create deephdc/deep-oc-obj_detect
(udocker) alvaro:~ $ udocker run -p 5000:5000 deephdc/deep-oc-obj_detect
Warning: this container exposes TCP/IP ports
Warning: non-existing user will be created
```

```
*****
*
*      STARTING 9bb8179d-c731-348c-bbf4-e43045ee89
*
*****
executing: deepaas-run
```

```
##      ###
##      ##### ##
.#####. #####. .#####.
##  ## ## //  ## //  ##  ##  ##
##. .## ### ### // ###  ##  ##
##  ##  ####  ####  #####.
Hybrid-DataCloud  ##
```

Welcome to the DEEPaaS API endpoint. You can directly browse the API documentation endpoint to check the API using the built-in Swagger UI or you can use any of our endpoints.

API documentation: <http://0.0.0.0:5000/ui>
API specification: <http://0.0.0.0:5000/swagger.json>
V2 endpoint: <http://0.0.0.0:5000/v2>




Swagger
Supported by SMARTBEAR

/swagger.json Explore

DEEP as a Service API endpoint ^{1.3.0}

/swagger.json



This is a REST API that is focused on providing access to machine learning models. By using the DEEPaaS API users can easily run a REST API in front of their model, thus accessing its functionality via HTTP calls.

Currently you are browsing the [Swagger UI](#) for this API, a tool that allows you to visualize and interact with the API and the underlying model.

- [Project website.](#)
- [Project documentation.](#)
- [Model marketplace.](#)

API documentation

versions

- GET** /v2/ Get V2 API version information
- GET** / Get available API versions

debug

- GET** /v2/debug/ Return debug information if enabled by API.

models

- GET** /v2/models/ Return loaded models and its information
- GET** /v2/models/obj_detect_pytorch/ Return model's metadata

Local inference and prediction

```
(udocker) alvaro:~ $ udocker pull deephdc/deep-oc-obj_detect
Info: downloading layer sha256:2746a4a261c9e18bfd7ff0429c18f
```


```
(...)
(udocker)
(udocker)
Warning
Warning
```

```
*****
*
*
*
*****
execut
```

```
.#
##
##.
#
```

```
Welcome
API doc
or you
```

```
API
API
```



V2 endpoint: `http://0.0.0.0:5000/v2`



/swagger.json

Explore

Service API endpoint ^{1.3.0}

providing access to machine learning models. By using the DEEPaaS API users can easily run a REST API in front of onality via HTTP calls.

[Swagger UI](#) for this API, a tool that allows you to visualize and interact with the API and the underlying model.

h information

ions

debug information if enabled by API.

GET /v2/models/ Return loaded models and its information

GET /v2/models/obj_detect_pytorch/ Return model's metadata

DEEPaaS: DEEP as a Service

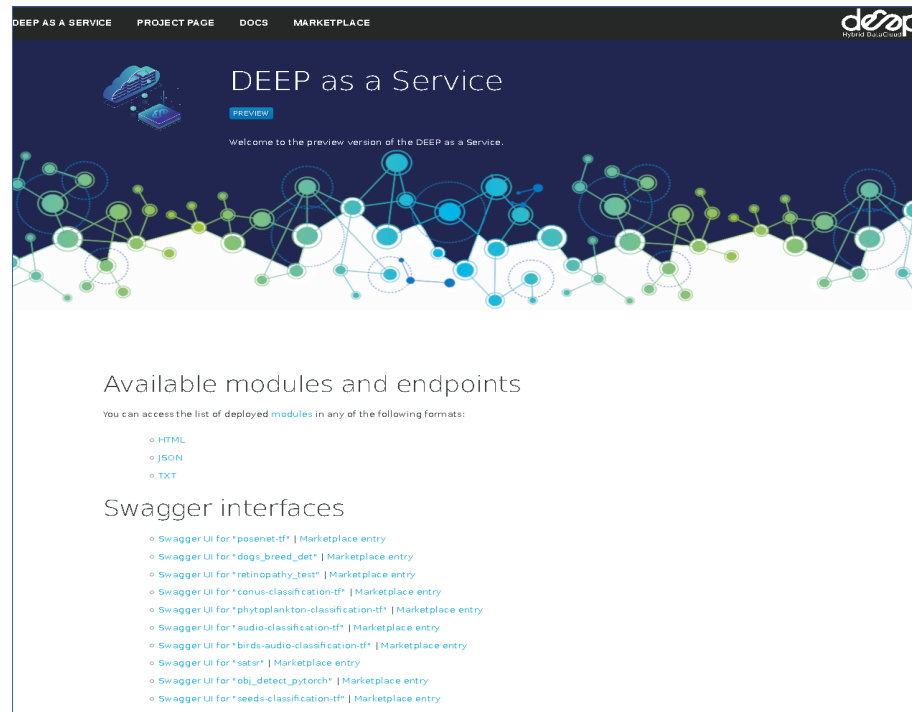
Automatic deployment of catalog modules to exploit their functionality

<https://deepaas.deep-hybrid-datacloud.eu>



Apache
MESOS™

APACHE
OpenWhisk™



DEEPaaS: automated deployment



The image shows a screenshot of the DEEP as a Service production endpoint interface. At the top left, there is a small icon of a cloud with a server rack and a cube with the DEEP logo. The main heading is "DEEP as a Service" in a large, white, sans-serif font. Below it, a yellow "PRODUCTION" label is visible. A welcome message reads "Welcome to the DEEP as a Service production endpoint." The background features a dark blue gradient with a white, jagged, mountain-like silhouette at the bottom. Overlaid on this is a complex network diagram with blue and green nodes connected by lines. Below the network diagram, the text "Available modules and endpoints" is displayed in a large, white, sans-serif font. Underneath this, a smaller line of text states: "You can access the list of deployed modules in any of the following formats:". This is followed by a bulleted list:

- [HTML](#)
- [JSON](#)
- [TXT](#)

 Below the list, the text "Swagger interfaces" is displayed in a large, white, sans-serif font. Underneath this, another bulleted list is shown:

- [Swagger UI for "image-classification-tf-dicom" | Marketplace entry](#)
- [Swagger UI for "dogs_breed_det" | Marketplace entry](#)
- [Swagger UI for "semseg_vaihingen" | Marketplace entry](#)
- [Swagger UI for "obj_detect_pytorch" | Marketplace entry](#)

Marketplace entry

DEEP OPEN CATALOG

PROJECT PAGE

DOCS

MODULES

CATEGORIES

deep
Hybrid DataCloud

Object Detection and Classification with Pytorch

A trained Region Convolutional Neural Network (Faster RCNN) for object detection and classification.

Model

Trainable

Inference

Pre-trained

Published by DEEP-Hybrid-DataCloud Consortium
Created: Thu 17 October 2019 - Updated: Mon 15 June 2020

Model Description

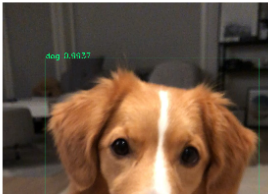
build

passing

This is a plug-and-play tool for object detection and classification using deep neural networks (Faster R-CNN ResNet-50 FPN Architecture [1]) that were already pretrained on the [COCO Dataset](#). The code uses the Pytorch Library, more information can be found at [Pytorch-Object-Detection](#).

The PREDICT method expects an image as input and will return a JSON with the predictions that are greater than the probability threshold. Let's say you have an image of a cat and a dog together and the probability output was 50% a dog and 80% a cat, if you set the threshold to 70%, the only detected object will be the cat, because its probability is grater than 70%.

This module works on uploaded images and gives as output the rectangle coordinates x1,y1 and x2,y2 were the classified object is located. It also provides you the probability of the classified detected object.



Categories

services, pytorch, rcnn, faster-rcnn, object detection, trainable, pre-trained, api-v2

License

License: MIT

Try it live! beta

WEB UI

API

API SPECIFICATION

CONFIGURE AND TRAIN

TRAINING DASHBOARD

Get the code

GITHUB

DOCKER HUB

Marketplace entry

DEEP OPEN CATALOGPROJECT PAGEDOCSMODULESCATEGORIES

Object Detection and Classification with Pytorch

A trained Region Convolutional Neural Network (Faster RCNN) for object detection and classification.

Model | Trainable | Inference | Pre-trained

deep Hybrid DataCloud

Categories

services, pytorch, rcnn, faster-rcnn, object detection, trainable, inference, pre-trained, api-v2

License

License: MIT

Try it live! beta

WEB UI

API

API SPECIFICATION

CONFIGURE AND TRAIN

TRAINING DASHBOARD

Get the code

GITHUB

DOCKER HUB



Swagger

https://deepaas.deep-hybrid-datacloud.eu/api/v1/web/deepaas/deep-oc/obj_detect_pytorch/swagger.json

Explore

DEEP as a Service API endpoint 1.3.0

[Base URL: /api/v1/web/deepaas/deep-oc/obj_detect_pytorch]
https://deepaas.deep-hybrid-datacloud.eu/api/v1/web/deepaas/deep-oc/obj_detect_pytorch/swagger.json


Hybrid DataCloud

This is a REST API that is focused on providing access to machine learning models. By using the DEEPaaS API users can easily run a REST API in front of their model, thus accessing its functionality via HTTP calls.

Currently you are browsing the [Swagger UI](#) for this API, a tool that allows you to visualize and interact with the API and the underlying model.

- Project website.
- Project documentation.
- Model marketplace.

API documentation

versions

GET /v2/ Get V2 API version information

GET / Get available API versions

debug

GET /v2/debug/ Return debug information if enabled by API.

models

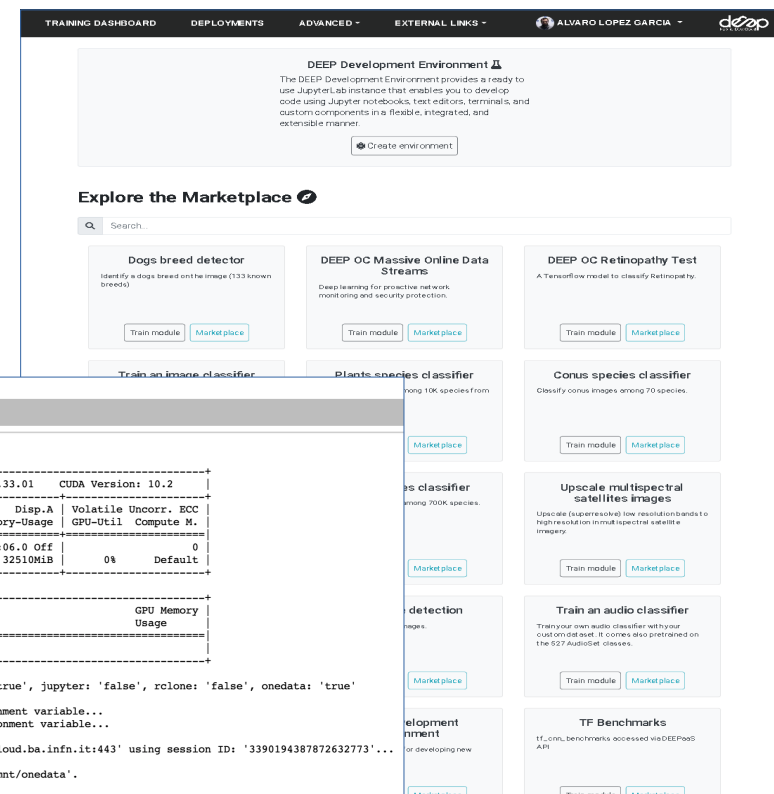
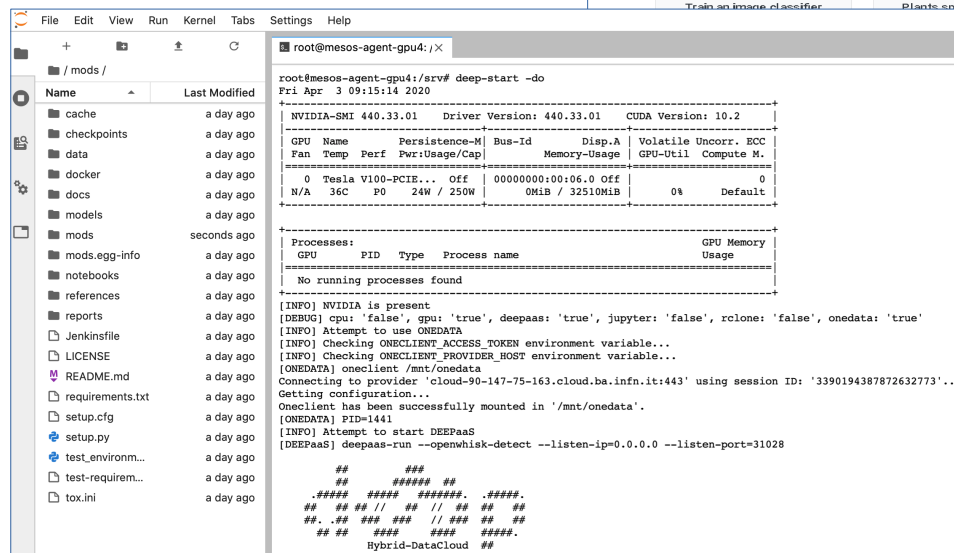
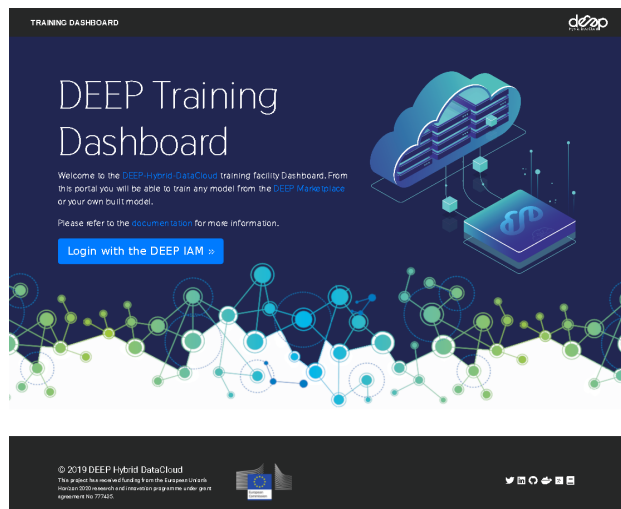
GET /v2/models/ Return loaded models and its information

GET /v2/models/obj_detect_pytorch/ Return model's metadata

DEEP training facility

1-click development and training environments in Cloud and HPC resources

<https://train.deep-hybrid-datacloud.eu>




DEEP training dashboard


TRAINING DASHBOARD


DEPLOYMENTS

ADVANCED ▾


EXTERNAL LINKS ▾


 ALVARO LOPEZ GARCIA ▾




DEEP Development Environment 

The DEEP Development Environment provides a ready to use JupyterLab instance that enables you to develop code using Jupyter notebooks, text editors, terminals, and custom components in a flexible, integrated, and extensible manner.

 [Create environment](#)

Explore the Marketplace 


 Search...

Dogs breed detector

Identify a dogs breed on the image (133 known breeds)

[Train module](#)

[Marketplace](#)



DEEP OC Massive Online Data Streams

Deep learning for proactive network monitoring and security protection.

[Train module](#)

[Marketplace](#)

DEEP OC Retinopathy Test

A Tensorflow model to classify Retinopathy.

[Train module](#)

[Marketplace](#)

Train an image classifier

Train your own image classifier with your custom dataset. It comes also pretrained on the 1K ImageNet classes.

[Train module](#)

[Marketplace](#)

Plants species classifier

Classify plant images among 10K species from the iNaturalist dataset.

[Train module](#)

[Marketplace](#)

Conus species classifier

Classify conus images among 70 species.

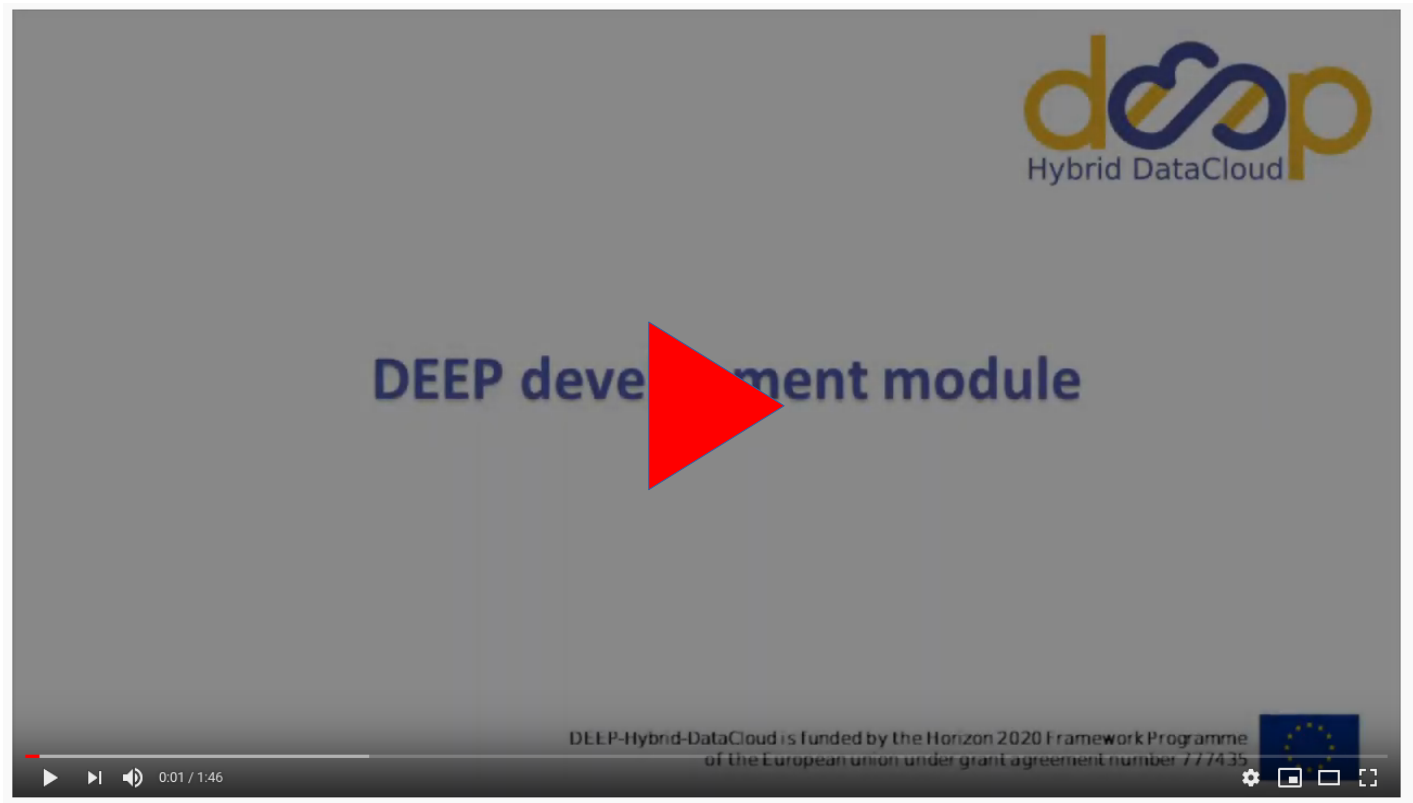
[Train module](#)

[Marketplace](#)

https://deep-hybrid-datacloud.eu

24/30

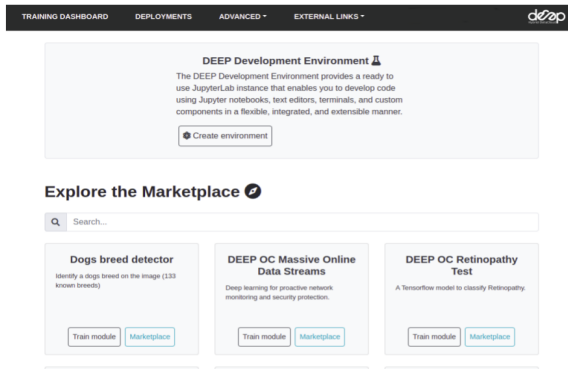
DEEP development environment



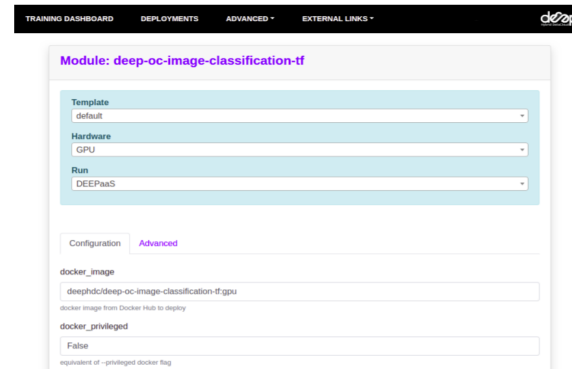
https://youtu.be/J_I_xWiBGNA

DEEP training facility


1. Select a module from the Open Catalogue



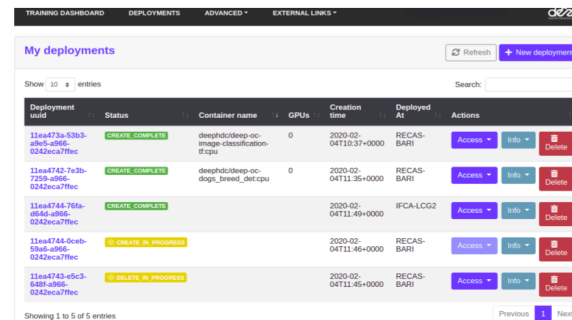
2. Choose your configuration parameters



4. Monitor the training history

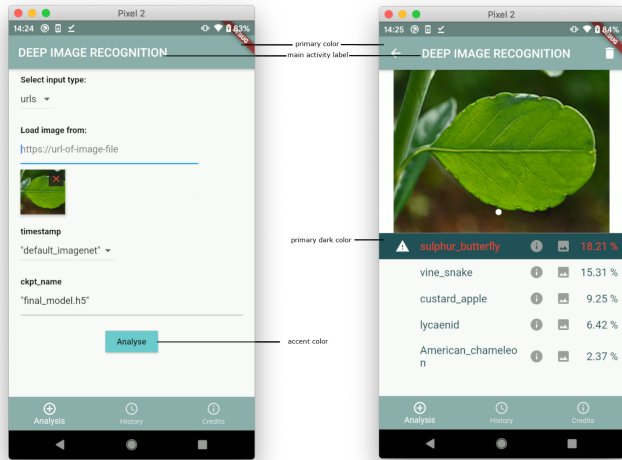


3. View your deployments



DEEPaaS API: models as services

Exposing models through a
common and standard-based API
<https://deepaas.readthedocs.io>



Deep Learning API

INTRO IMAGE RECOGNITION MISCELLANEOUS

This webpage gathers all the applications of the Instituto de Física de Cantabria (IFCA) using deep learning techniques.

Image Recognition

Play

Author

Description

This application classifies a plant image among 6K plant species (mainly from Western Europe).

DEEP as a Service API endpoint

1.0.0.dev32

Hybrid DataCloud

This is a REST API that is focused on providing access to machine learning models. By using the DEEPaaS API users can easily run a REST API in front of their model, thus accessing its functionality via HTTP calls.

Currently you are browsing the [Swagger UI](#) for this API, a tool that allows you to visualize and interact with the API and the underlying model.

- Project website.
- Project documentation.
- Model marketplace.

Contact the developer:
Apache 2.0
API documentation

versions

- HEAD /v2/ Get V2 API version information
- GET /v2/ Get V2 API version information
- GET / Get available API versions

debug

- HEAD /v2/debug Return debug information if enabled by API.
- GET /v2/debug Return debug information if enabled by API.

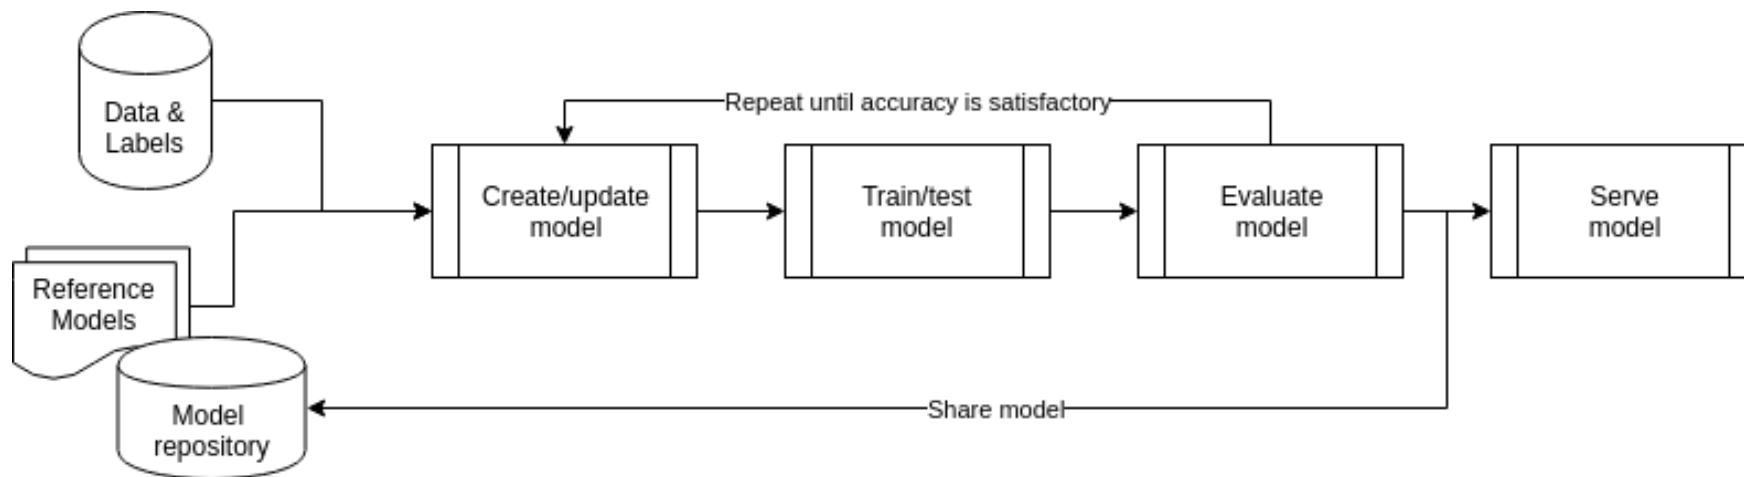
models

- HEAD /v2/models Return loaded models and its information
- GET /v2/models Return loaded models and its information
- HEAD /v2/models/deepaas-test Return model's metadata
- GET /v2/models/deepaas-test Return model's metadata
- POST /v2/models/deepaas-test/train Retrain model with available data
- HEAD /v2/models/deepaas-test/train Get a list of trainings (running or completed)
- GET /v2/models/deepaas-test/train Get a list of trainings (running or completed)
- HEAD /v2/models/deepaas-test/train/{uuid} Get status of a training
- GET /v2/models/deepaas-test/train/{uuid} Get status of a training
- DELETE /v2/models/deepaas-test/train/{uuid} Cancel a running training
- POST /v2/models/deepaas-test/predict Make a prediction given the input data

Models

- Failure >
- Location >
- Version >

Machine learning development cycle



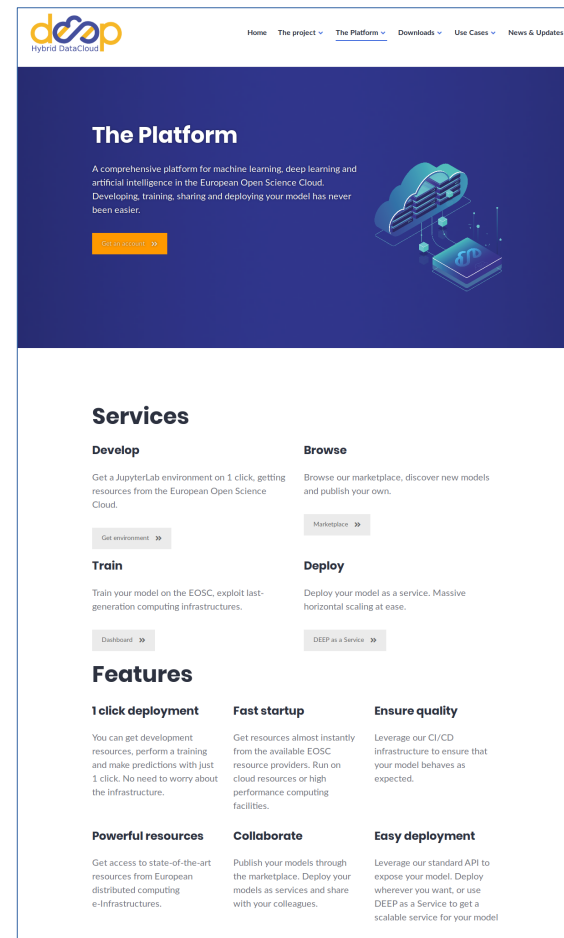
- The project covers **all development phases** of a machine learning application:
 - Creation, update, improvement of a model (existing or new) → **DEEP Open Catalog and DEEP training facility**
 - Model training, test, evaluation → **DEEP training facility**
 - Model deployment as a service → **DEEP as a Service**
 - Model publication for sharing and reuse → **DEEP Open Catalog**

DEEP-Rosetta: 2nd platform and release



Second software release and platform released January 2020

- Full support for the whole machine learning application life cycle
 - Sharing and reuse through DEEP Open Catalogue
 - Training, testing and evaluation on DEEP training facility
 - Live deployment on DEEPaaS
- Transparent access to Cloud and HPC resources



Conclusions

- DEEP provides a comprehensive, service-oriented, framework to develop machine learning and artificial intelligence applications
 - Exploiting distributed e-Infrastructure and HPC resources, with special focus on the EOSC
 - Covering the whole development phases (CRISP-DM and machine learning lifecycle)
 - High level and added value services that can be used standalone or as a whole platform solution
- Services at production level with TRL \geq 8
- Industry opportunities: engagement of industrial stakeholders through collaboration with EOSC-Digital Innovation Hub

A Cloud-Based Framework for Machine Learning Workloads and Applications

ÁLVARO LÓPEZ GARCÍA¹, JESÚS MARCO DE LUCAS¹, MARICA ANTONACCI², WOLFGANG ZU CASTELL^{3,10}, MARIO DAVID¹², MARCUS PŁOCIENNIK¹³, VIET TRAN¹⁴, LARA LLORET IGLESÍAS¹⁵, GERMÁN MOLTÓ¹⁶, MARCIN PŁOCIENNIK¹⁷, ANDY S. ALIC¹⁸, MIGUEL CABALLER¹⁹, ISABEL CAMPOS PLASENCIA²⁰, ALESSANDRO COSTANTINI²¹, STEFAN DŁUGOLINSKY²², DOINA CRISTINA DUMA²³, GIACINTO DONVITO²⁴, JORGE GOMES²⁵, IGNACIO HEREDIA CACHA²⁶, KEIICHI ITO²⁷, VALENTIN Y. KOZLOV²⁸, GIANG NGUYEN²⁹, PABLO ORVIZ FERNÁNDEZ³⁰, ZDENĚK ŠUST³¹, AND PAWEŁ WOLNIEWICZ³²

¹ICAC3-UC3, 39003 Santander, Spain

²Laboratory of Instrumentation and Experimental Particle Physics, 1000-078 Lisbon, Portugal

³IN3N Bari, 70126 Bari, Italy

⁴IN3N CNR, 40127 Bologna, Italy

⁵Poznań Supercomputing and Networking Center, IBCh PAS, 61-704 Poznań, Poland

⁶Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

⁷Instituto de Instrumentación para Imagen Molecular (IIM), CSIC, Universidad Politécnica de Valencia, 46102 Valencia, Spain

⁸Institute of Informatics, Slovak Academy of Sciences (IAS), 845 07 Bratislava, Slovakia

⁹Leibniz-Zentrum München, Deutsches Forschungszentrum für Gesundheit, 85748 Oberschleißheim, Germany

¹⁰Department of Mathematics, Technische Universität München, 85748 Munich, Germany

¹¹CEISNET, 16000 Prague, Czech Republic

Corresponding author: Álvaro López García (alopez@icac3.com)

This work was supported by the project DEEP-Hybrid DataCloud “Designing and Enabling E-infrastructures for intensive Processing in a Hybrid DataCloud” that has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under Grant 777435.

ABSTRACT In this paper we propose a distributed architecture to provide machine learning practitioners with a set of tools and cloud services that cover the whole machine learning development cycle: ranging from the models creation, training, validation and testing to the models serving as a service, sharing and publication. In such respect, the DEEP-Hybrid-DataCloud framework allows transparent access to existing e-Infrastructures, effectively exploiting distributed resources for the most compute-intensive tasks coming from the machine learning development cycle. Moreover, it provides scientists with a set of Cloud-oriented services to make their models publicly available, by adopting a serverless architecture and a DevOps approach, allowing an easy share, publish and deploy of the developed models.

INDEX TERMS Cloud computing, computers and information processing, deep learning, distributed computing, machine learning, serverless architectures.

1. INTRODUCTION

The impact of emerging computing techniques together with an increasing dimension of large datasets and the availability of more and more performing and accessible computing resources is transforming many research areas. This opens the door to new opportunities to tackle unprecedented research challenges. Over the last decade there has been a boost on the usage of machine learning techniques in most of the research

areas, and recently it has even improved with the adoption of deep learning techniques, e.g. LeCun et al. [1]. Although the basic components of the techniques are well known, recent advances arouse the interest from the scientific community towards this area, and it has already become a state-of-the-art technology in many fields, from computer vision to speech recognition.

The performance increase in the existing computing technologies and the availability of specialized computing devices played an important role in the advent of deep learning. In particular, the availability of more efficient and

The associate editor coordinating the review of this manuscript and approving it for publication was Liangxin Hu.

Request access through EOSC portal

<https://marketplace.eosc-portal.eu/services/deepaas-training-facility>

DEEP Services

Marketplace: <https://marketplace.deep-hybrid-datacloud.eu>

Training facility: <https://train.deep-hybrid-datacloud.eu/>

DEEPaaS: <https://deepaas.deep-hybrid-datacloud.eu/>

Documentation

<http://docs.deep-hybrid-datacloud.eu>

DEEP videos

https://www.youtube.com/playlist?list=PLJ9x9Zk1O-J_UZfNO2uWp2pFMmbwLvzXa

Contact

aloga@ifca.unican.es, deep-info@listas.csic.es

Web page

<https://deep-hybrid-datacloud.eu/the-platform/>

Social media

https://twitter.com/DEEP_eu

https://twitter.com/IFCA_Computing

Thank you
Any Questions?



<https://deep-hybrid-datacloud.eu>



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 777435.