Lecture 2: Virtual Machines & Virtual Environments

AC215

Pavlos Protopapas
SEAS/Harvard

Outline

1. Motivation
2. Virtual Machines
3. Virtual Environments

Outline

1. **Motivation**
2. Virtual Machines
3. Virtual Environments

Motivation: Deep Learning Flow

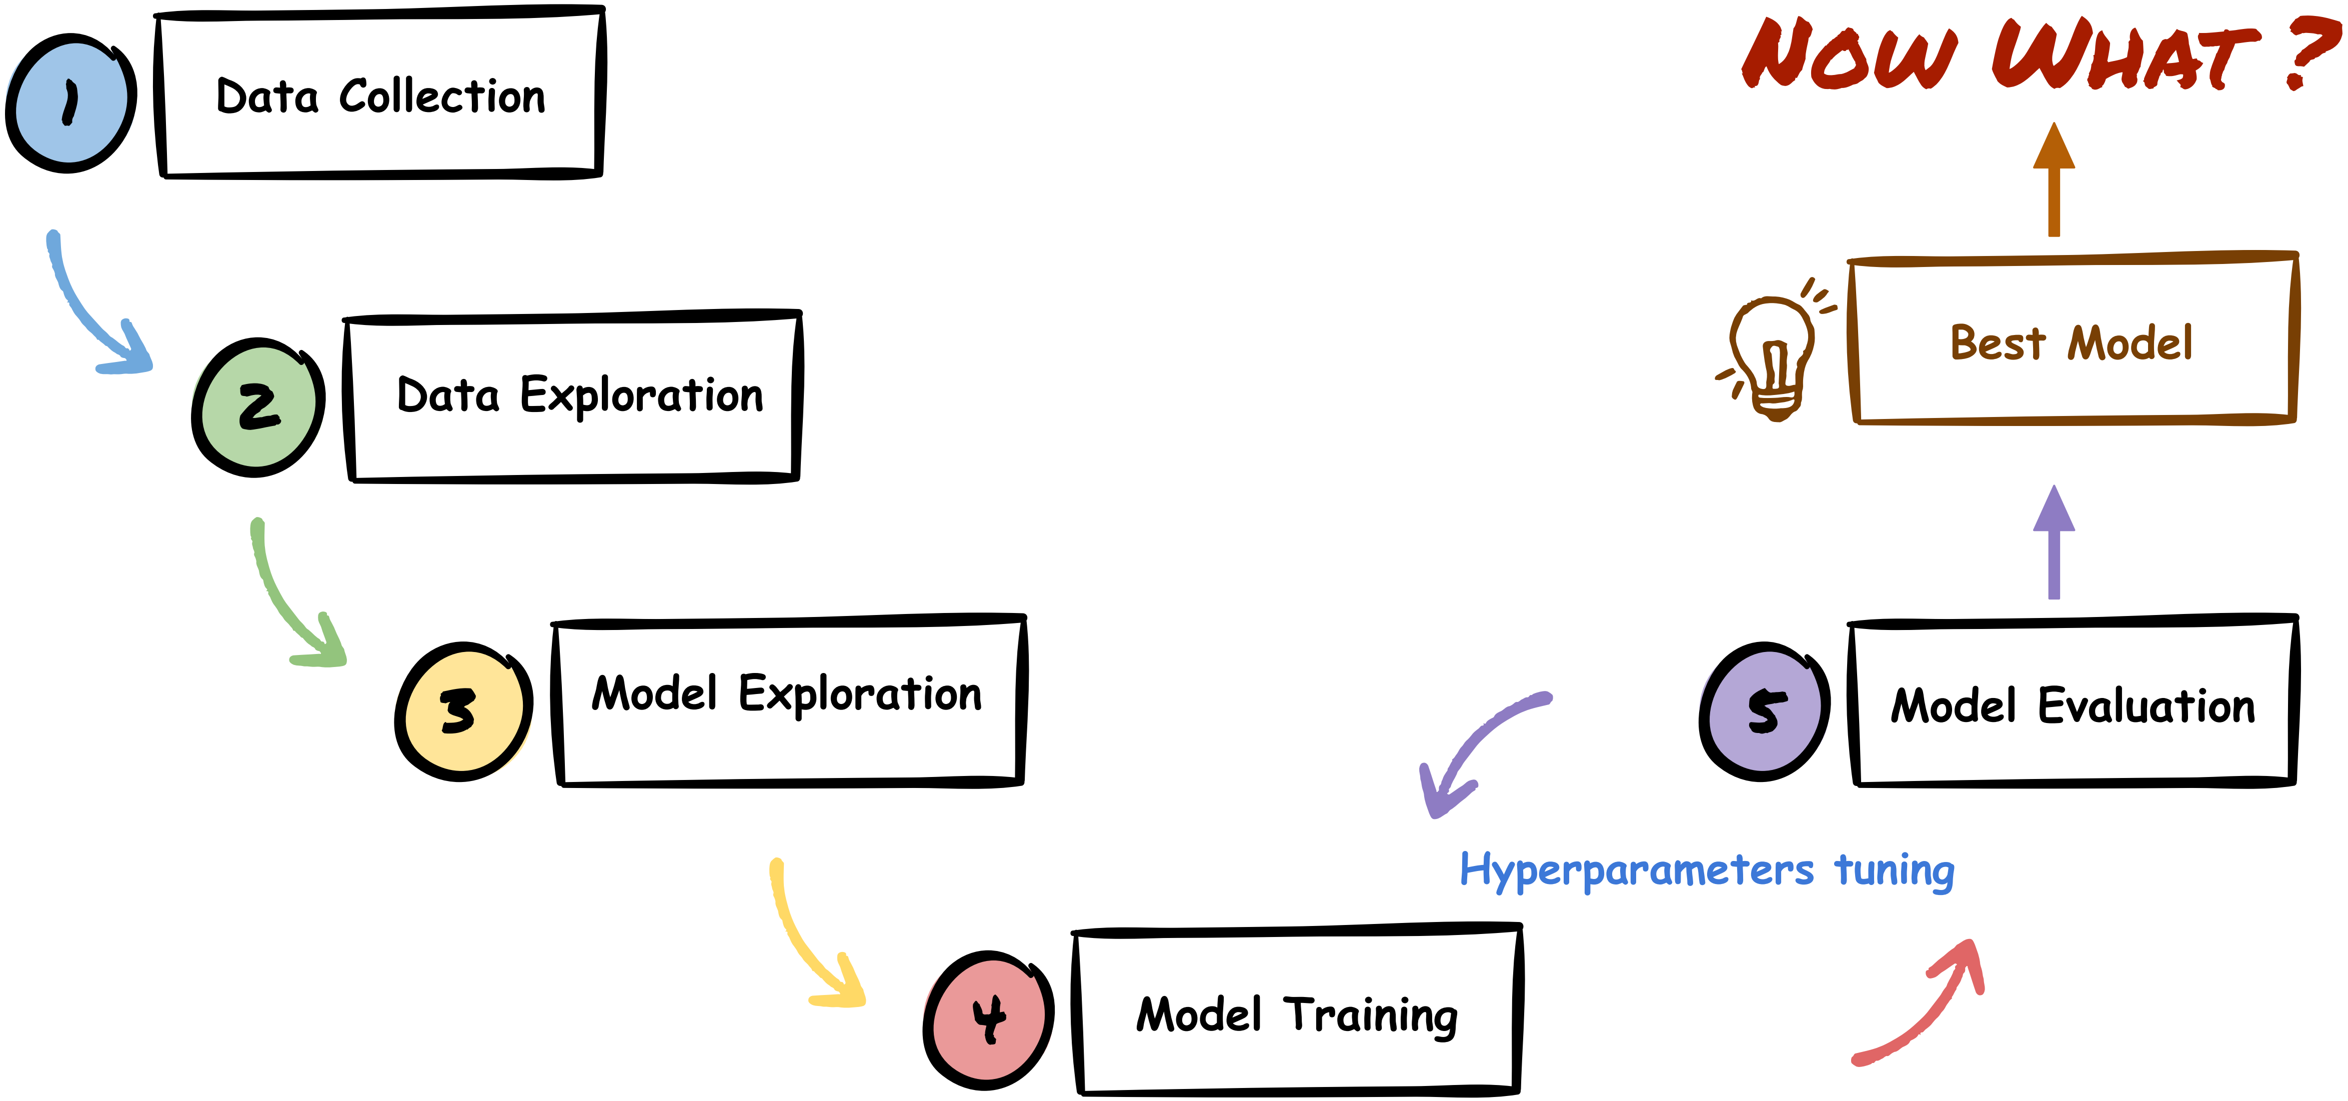

Motivation: Best Model

trainable_parameters	execution_time	loss	accuracy	model_size	learning_rate	batch_size	epochs	optimizer	name
2,306,051	2.97 mins	42.87	90.91%	10 MB	0.001	32	10	SGD	tfhub_mobilenetv2_train_base_True
82,179	3.19 mins	42.79	90.30%	10 MB	0.001	32	10	SGD	tfhub_mobilenetv2_train_base_False
164,355	3.91 mins	70.97	89.09%	10 MB	0.001	32	15	SGD	mobilenetv2_train_base_False
2,388,227	2.95 mins	82.03	88.48%	10 MB	0.001	32	10	SGD	mobilenetv2_train_base_True
11,112,323	6.85 mins	0.79	67.88%	44 MB	0.010	32	25	SGD	4_block
25,950,531	8.19 mins	0.74	66.67%	104 MB	0.010	32	25	SGD	2_block
22,514,755	4.78 mins	1.07	41.21%	90 MB	0.010	32	15	SGD	vgg_style

We want to build a 🍄 Mushroom Finder App

- Pavlos likes to go the forest for mushroom picking
- Some mushrooms can be poisonous
- Help build an app to identify mushroom type and if poisonous or not

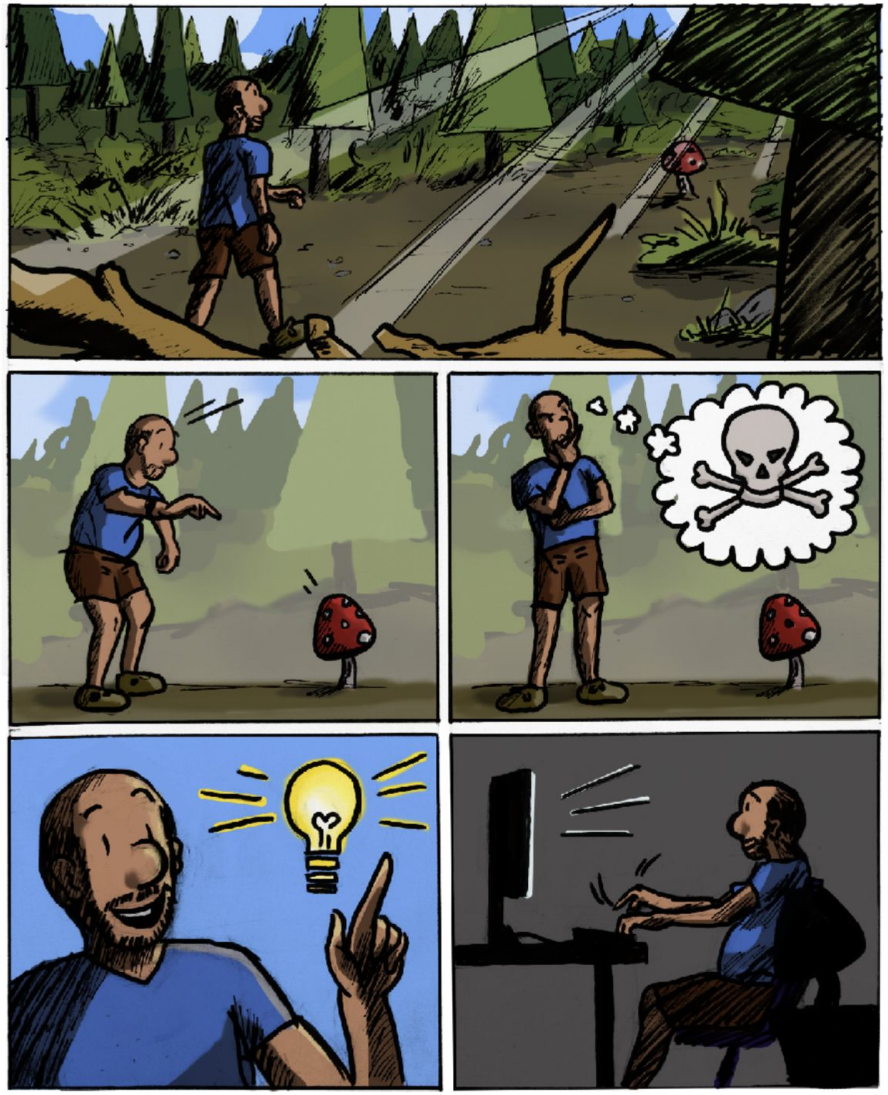

Credit: Nikolas Protopapas

🍄 Mushroom App: Data

- Collect images from Google
- For our demo we downloaded images for mushrooms oyster, crimini, **amanita (Poisonous)**
- Images organized into 3 labels

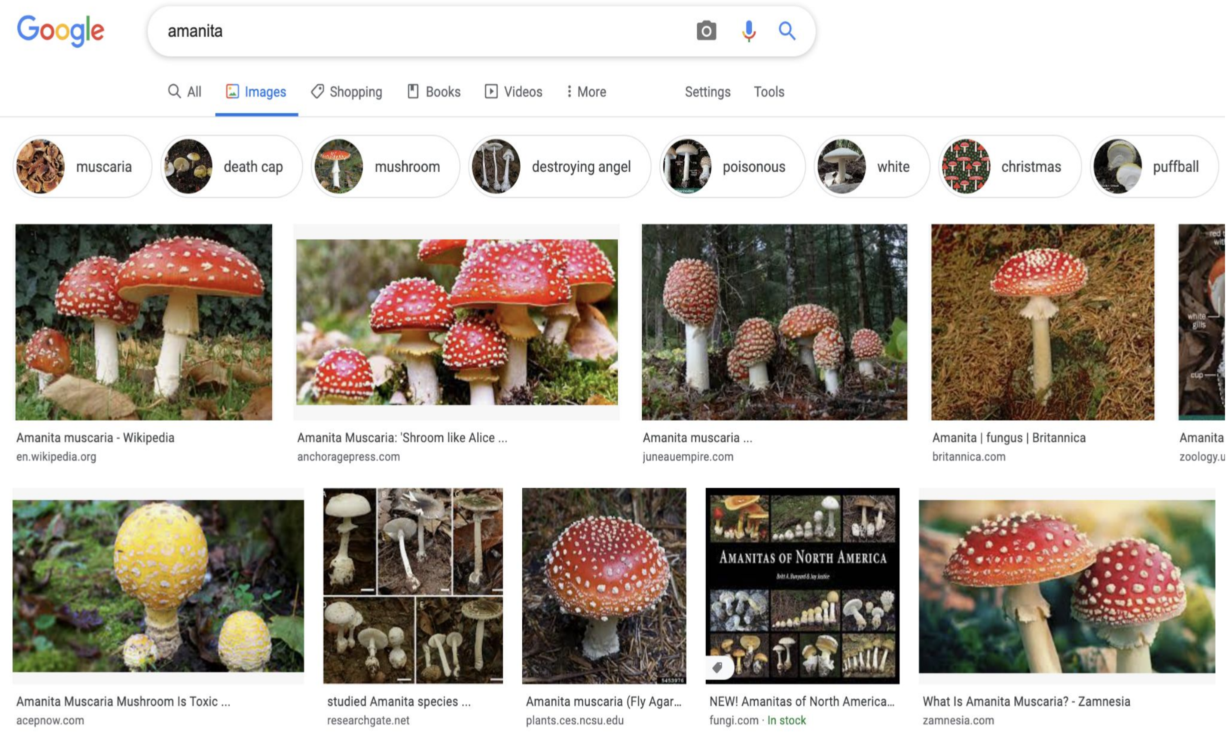

Python Script

🍄 Mushroom App: Models

- Identify our problem task
- Try various model architectures
- Transfer Learning
- Hyperparameters tuning
- Experiment Tracking

trainable_parameters	execution_time	loss	accuracy
2,306,051	2.97 mins	42.87	90.91%
82,179	3.19 mins	42.79	90.30%
164,355	3.91 mins	70.97	89.09%
2,388,227	2.95 mins	82.03	88.48%
11,112,323	6.85 mins	0.79	67.88%
25,950,531	8.19 mins	0.74	66.67%
22,514,755	4.78 mins	1.07	41.21%

Colab

🍄 Mushroom App: Best Model

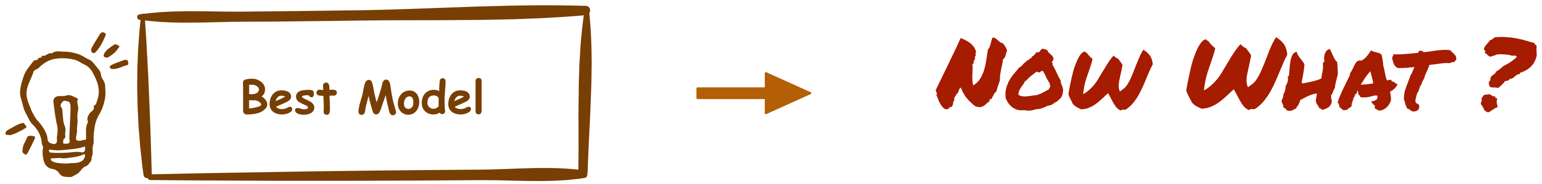

Mushroom App

- We want to build an app to take a photo of a mushroom and it helps us identify the type of mushroom
- How do we build the app?

How do we build an App?

- Collaborate with team to design and develop
- Expose best model as an API
- Build a frontend using HTML & javascript
- Integrate model prediction API into the app
- Deploy app to a cloud provider
- http://awesome-mushroom-app.com [Go live]

How do we build an App?

Data Collection

Data Exploration

Model Exploration

Model Training

Model Evaluation

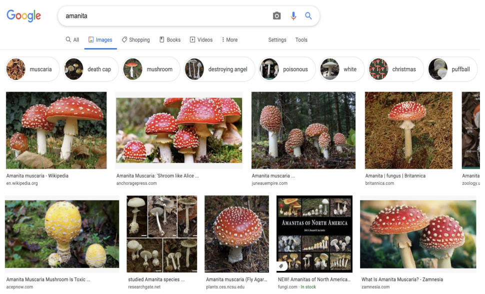

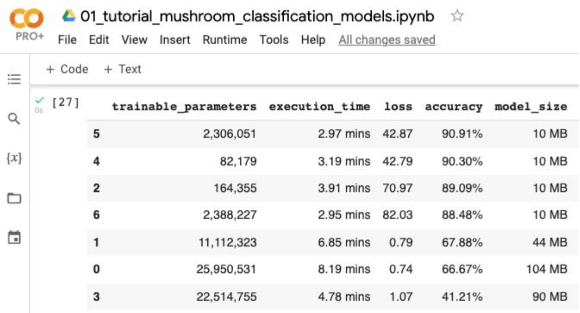

	trainable_parameters	execution_time	loss	accuracy	model_size
5	2,306,051	2.97 mins	42.87	90.91%	10 MB
4	82,179	3.19 mins	42.79	90.30%	10 MB
2	164,355	3.91 mins	70.97	89.09%	10 MB
6	2,388,227	2.95 mins	82.03	88.48%	10 MB
1	11,112,323	6.85 mins	0.79	67.88%	44 MB
0	25,950,531	8.19 mins	0.74	66.67%	104 MB
3	22,514,755	4.78 mins	1.07	41.21%	90 MB

Rest API

Best Model

Type: amanita (93.54%)

Python Script

Colab

IDE / Code Editor

How do we build an App?

Data Collection

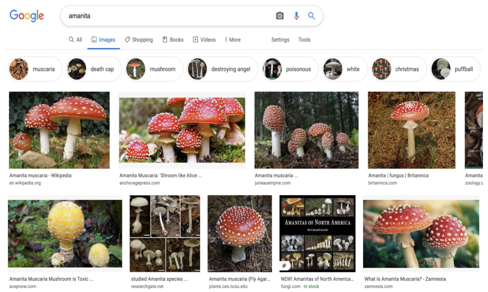

Python Script

Data Exploration

Model Exploration

Model Training

Model Evaluation

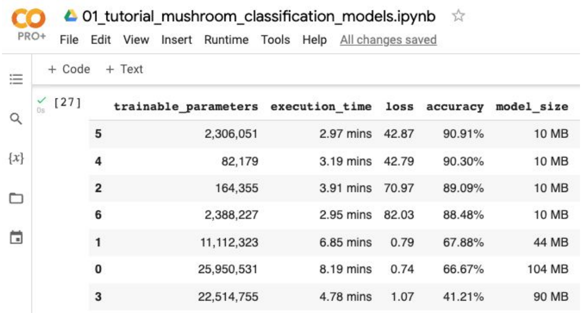

	trainable_parameters	execution_time	loss	accuracy	model_size
5	2,306,051	2.97 mins	42.87	90.91%	10 MB
4	82,179	3.19 mins	42.79	90.30%	10 MB
2	164,355	3.91 mins	70.97	89.09%	10 MB
6	2,388,227	2.95 mins	82.03	88.48%	10 MB
1	11,112,323	6.85 mins	0.79	67.88%	44 MB
0	25,950,531	8.19 mins	0.74	66.67%	104 MB
3	22,514,755	4.78 mins	1.07	41.21%	90 MB

Colab

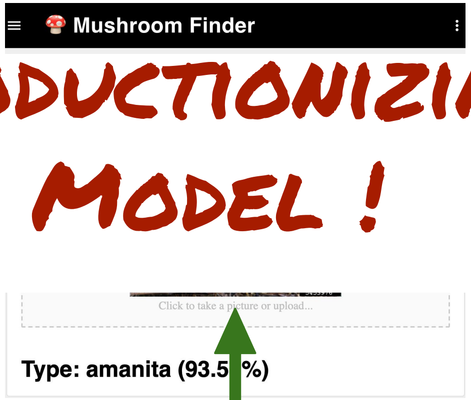

PRODUCTIONIZING MODEL !

Rest API

Best Model

Type: amanita (93.54%)

IDE / Code Editor

Challenges

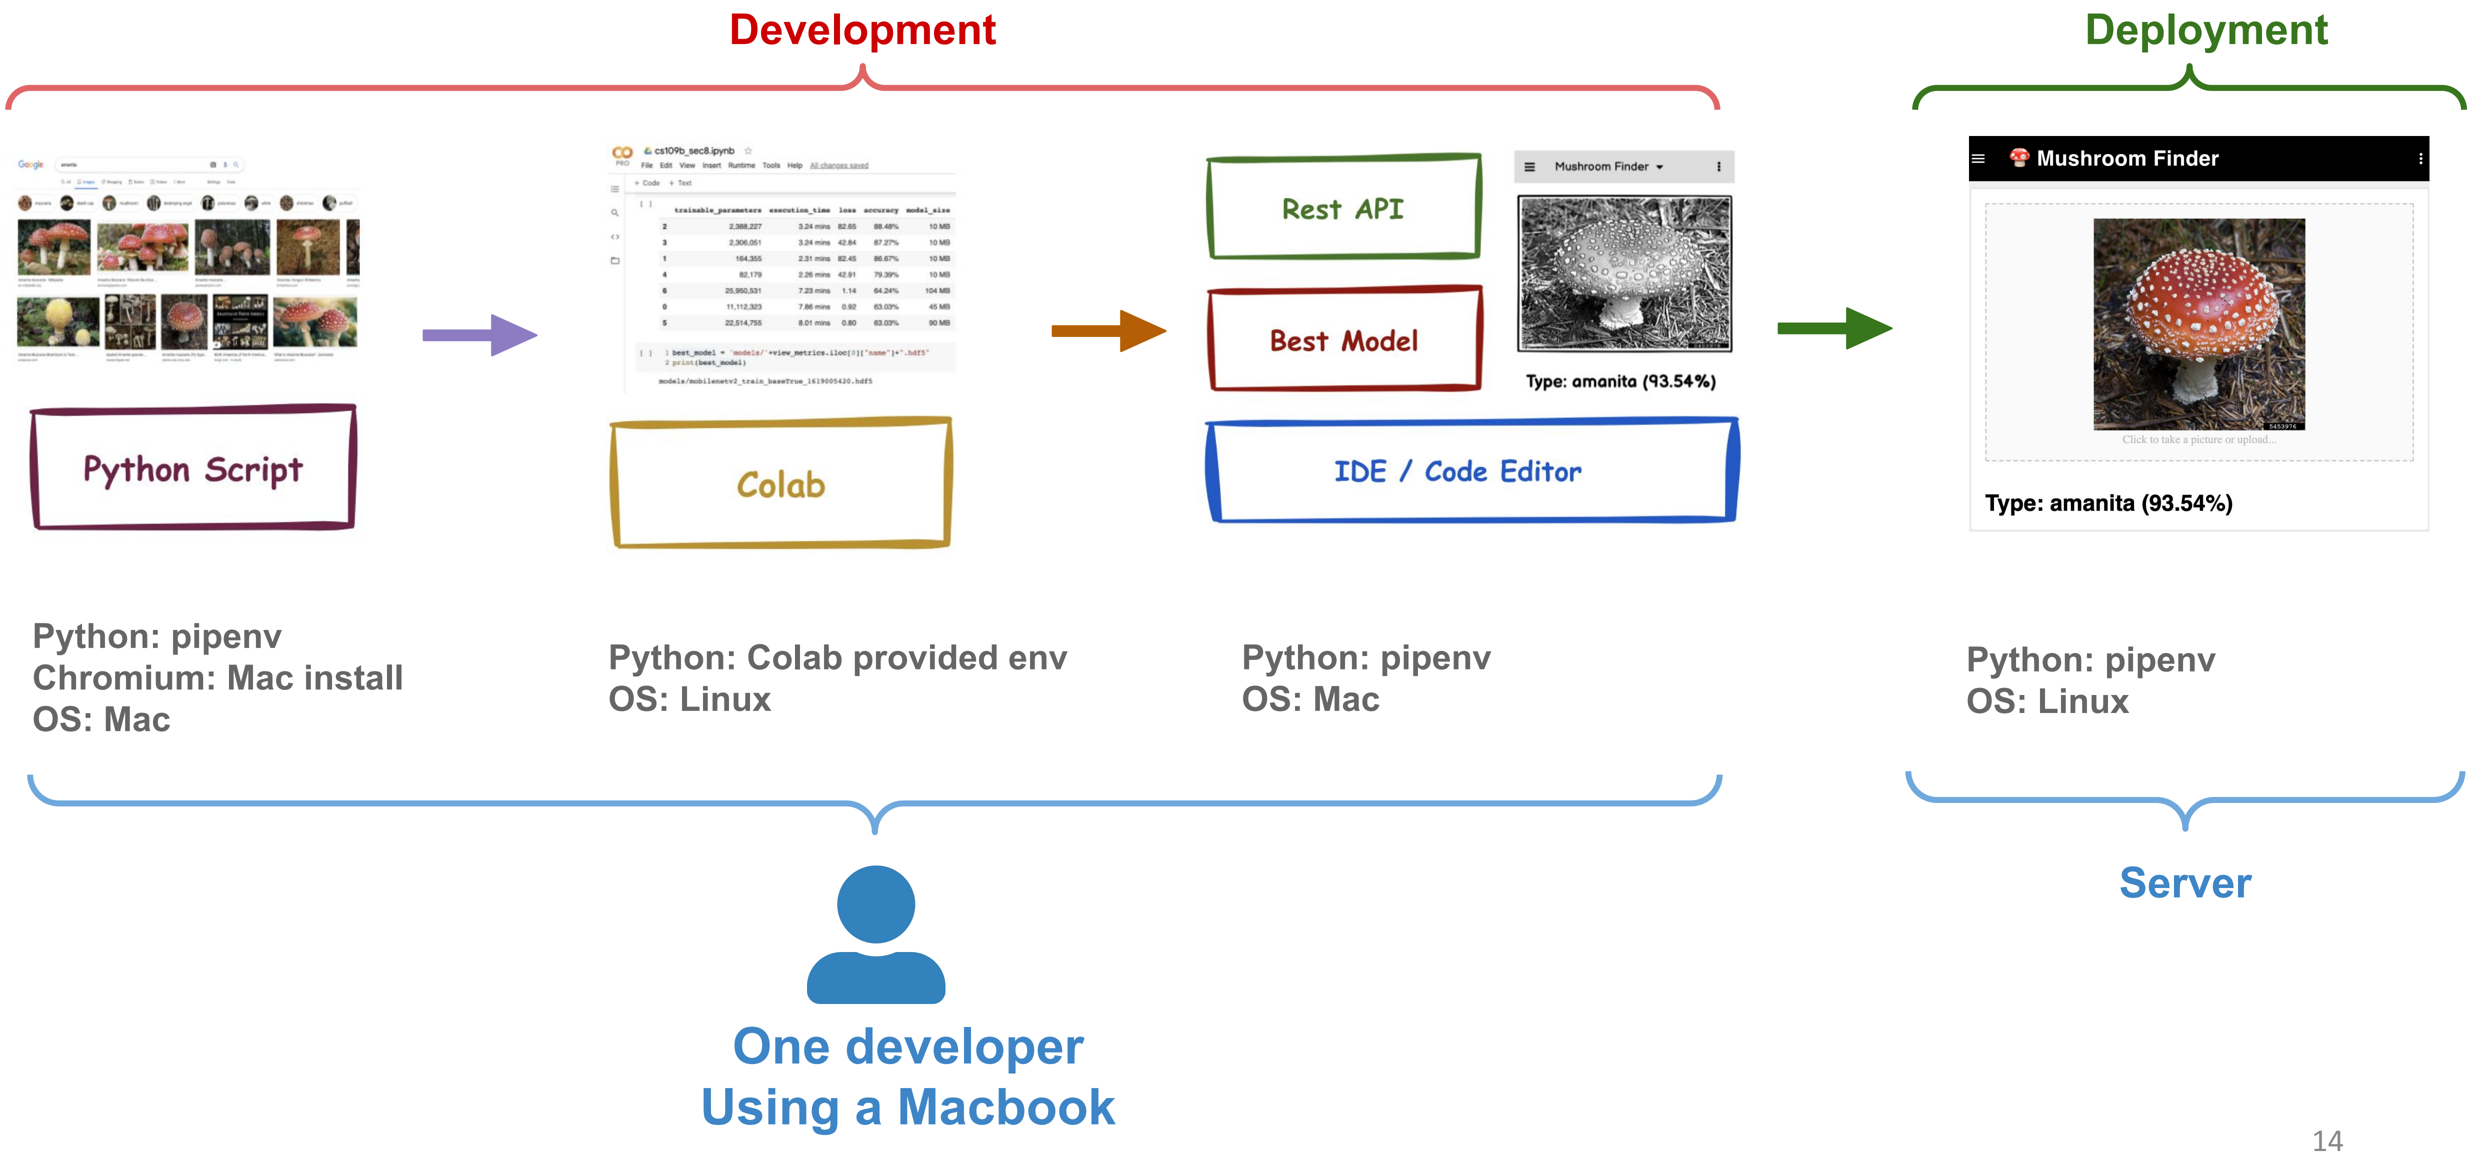

Development
Deployment
Python Script
Colab
Rest API
Best Model
Mushroom Finder
Type: amanita (93.54%)
IDE / Code Editor
Mushroom Finder
Type: amanita (93.54%)
Python: pipenv
Chromium: Mac install
OS: Mac
Python: Colab provided env
OS: Linux
Python: pipenv
OS: Mac
Python: pipenv
OS: Linux
Server
One developer
Using a Macbook

Challenges - Multiple Developers

Development

Deployment

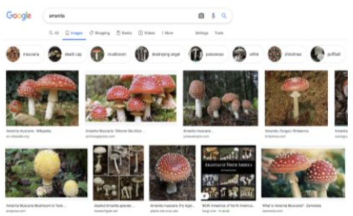

Python Script

Colab

Rest API

Best Model

IDE / Code Editor

Python: pipenv
Chromium: Mac install, Windows install
OS: Mac, Windows

Python: Colab provided env
OS: Linux

Python: pipenv
OS: Mac, Windows

Python: pipenv
OS: Linux

Server

Multiple developers, Using Mac and Windows OS

Challenges - Multiple Developers + Automation

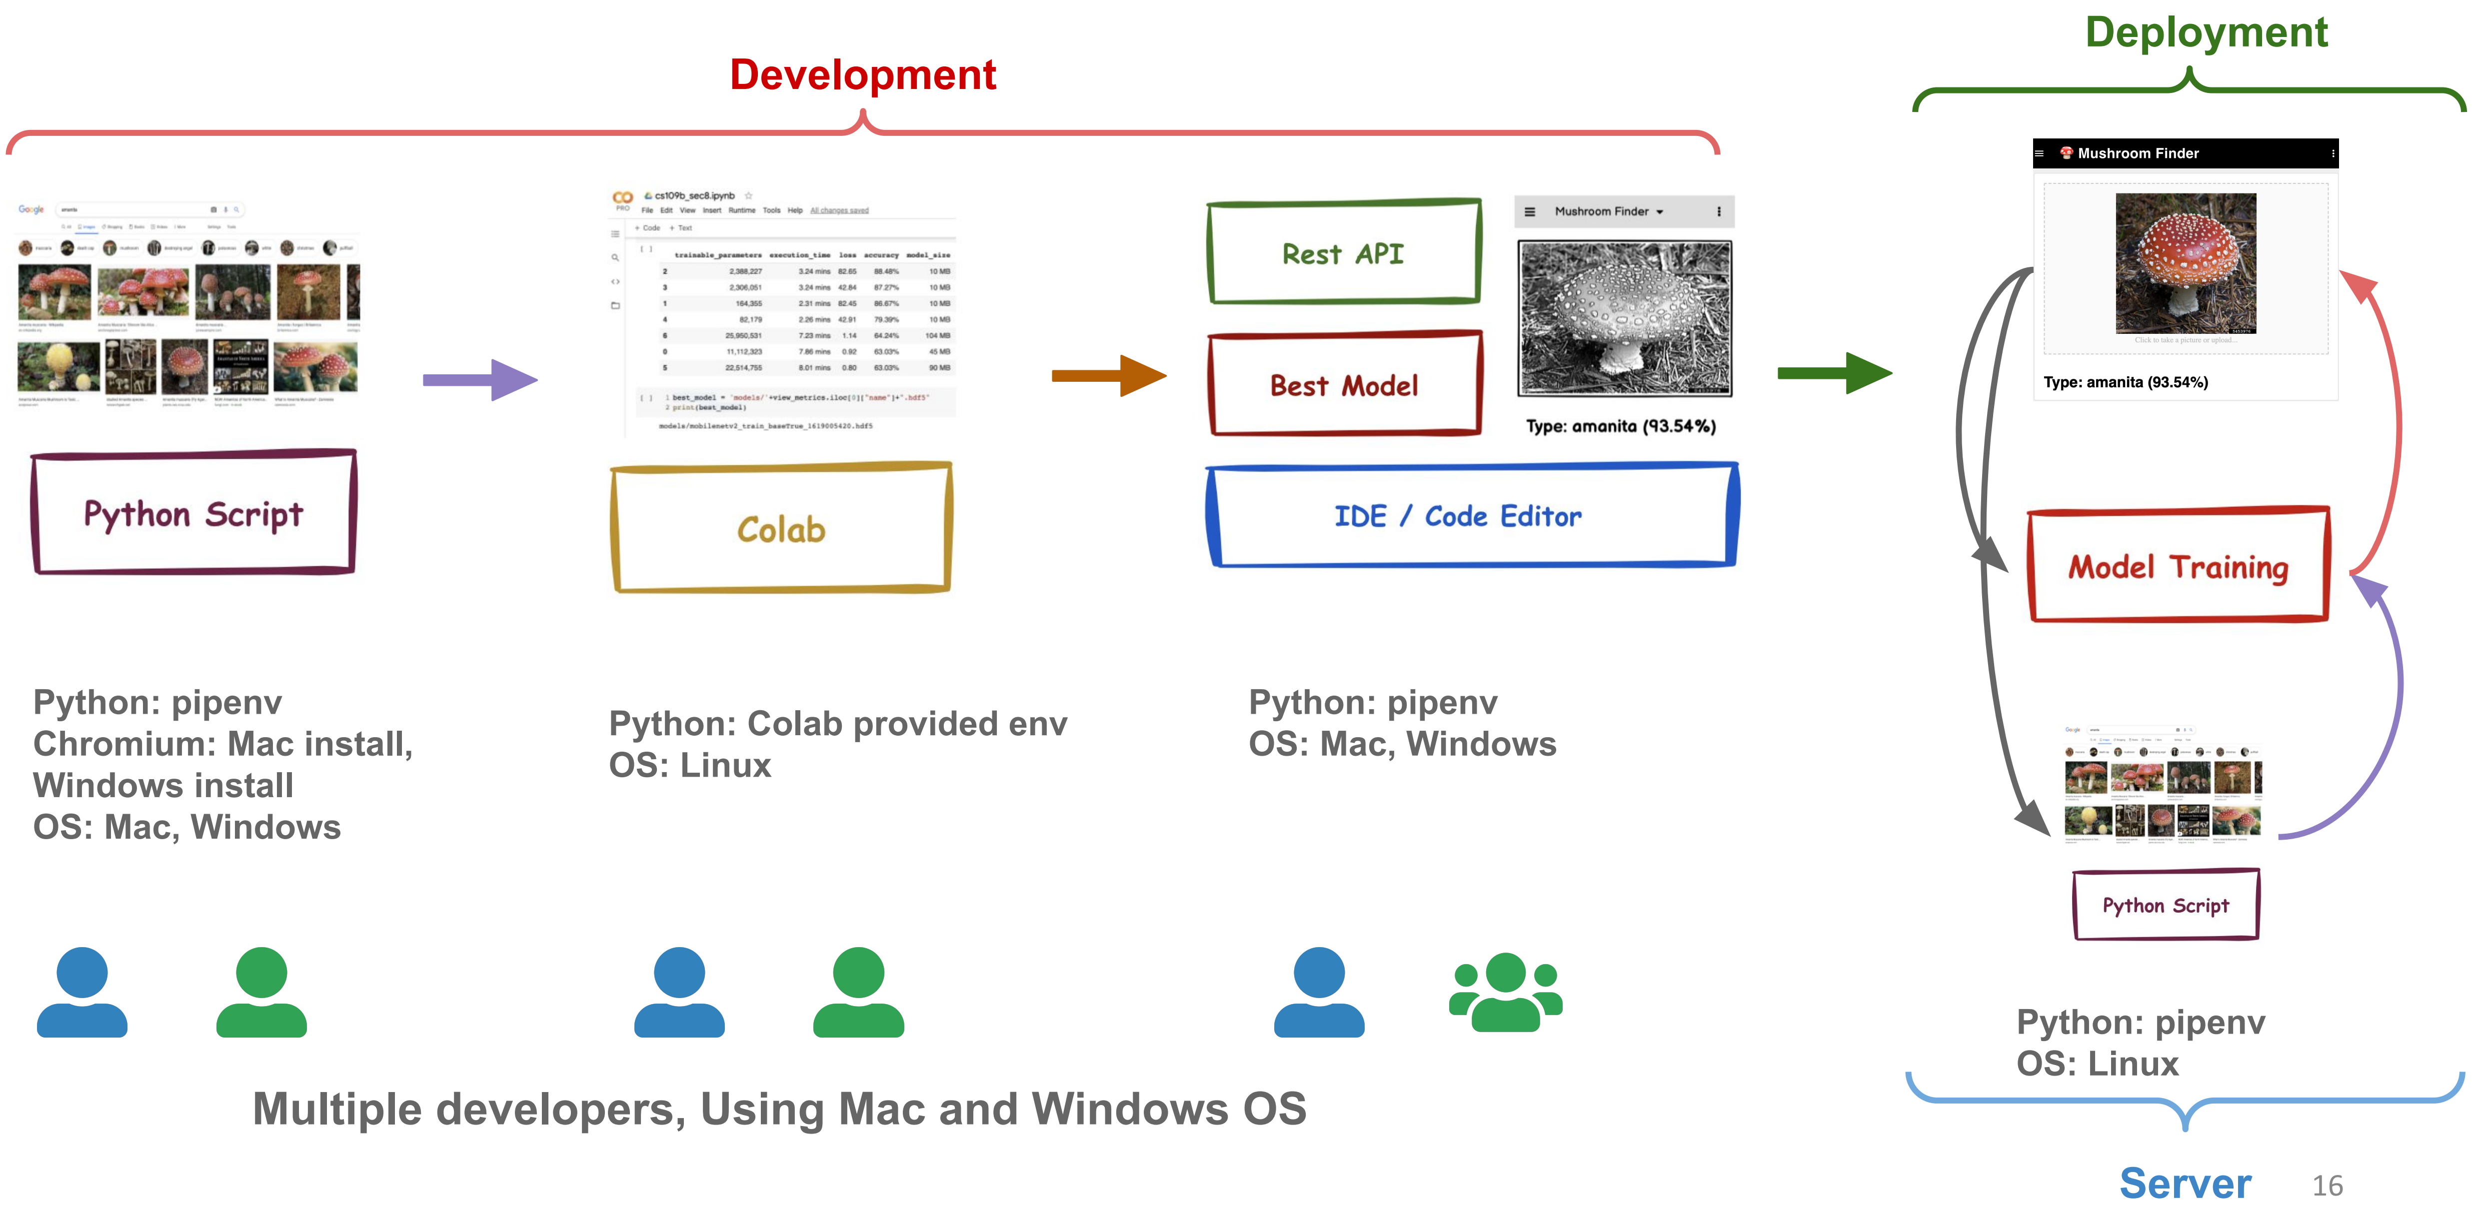

Challenges / Solutions

Challenges:

- Required Installations for Specific Operating Systems
- Guidelines for Code Collaboration
- Methods for Sharing Datasets and Models
- Automation of Data Gathering and Model Training
- Onboarding Procedures for New Team Members
- Resolving "It Works on My Machine" Issues ¯_(ツ)_/¯

Solutions:

- Isolate development into environments that can be shared
- Develop in a common OS regardless of developers host OS
- Track software/framework installs

Tools

- Virtual Machines
- Virtual Environments
- Containers

Outline

1. Motivation
2. **Virtual Machines**
3. Virtual Environments

Running the Simple-Translate App on a **Virtual Machine**
To achieve this, follow the steps below:

- Create a Virtual Machine Instance.
- SSH into the Virtual Machine.
- Install Required Dependencies: git, Python.
- Download and Execute the Simple-Translate Python Script.
- For detailed instructions, please refer to the following link: Installing App on VM Manually. (https://github.com/dlops-io/simple-translate#installing-app-on-vm-manually)

Virtual Machines Tutorial

Google Cloud Platform: https://cloud.google.com

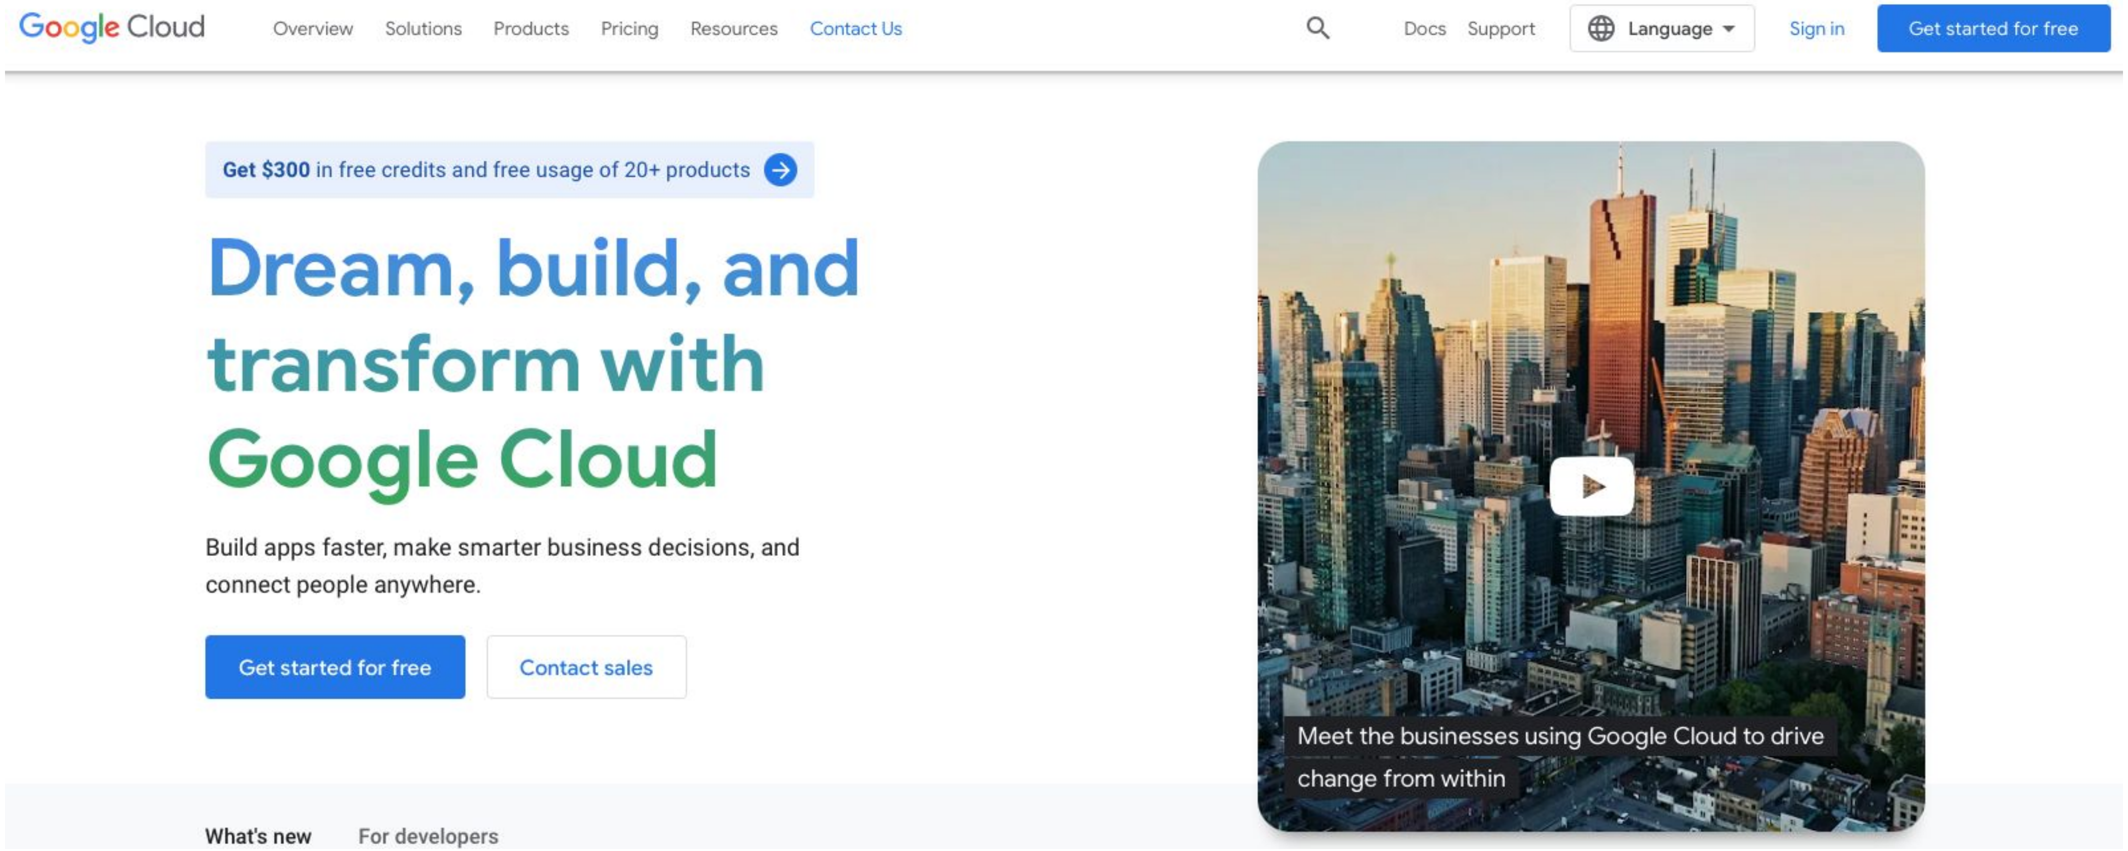

Virtual Machines Tutorial

Go to Navigation Menu

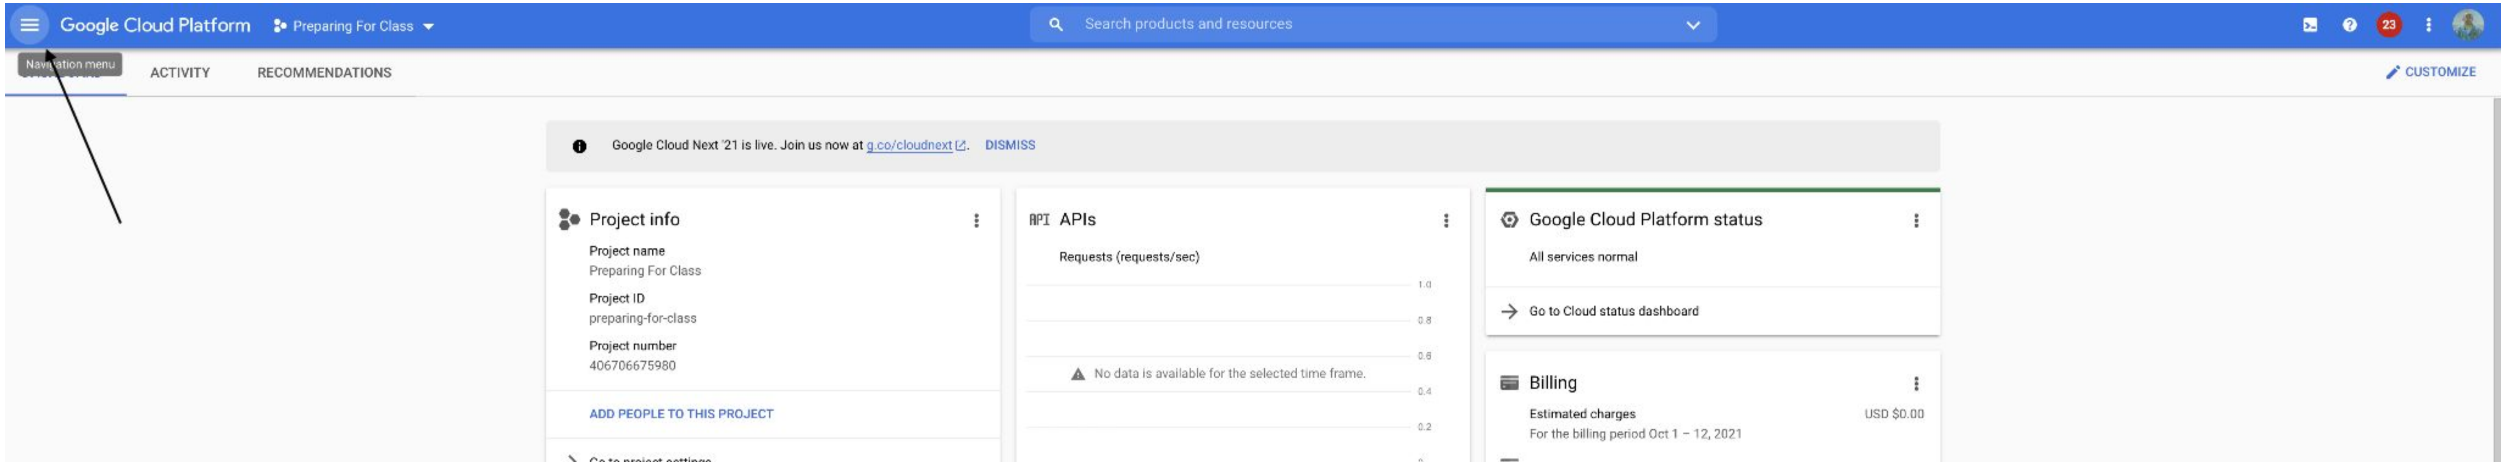

Virtual Machines Tutorial

Select compute engine

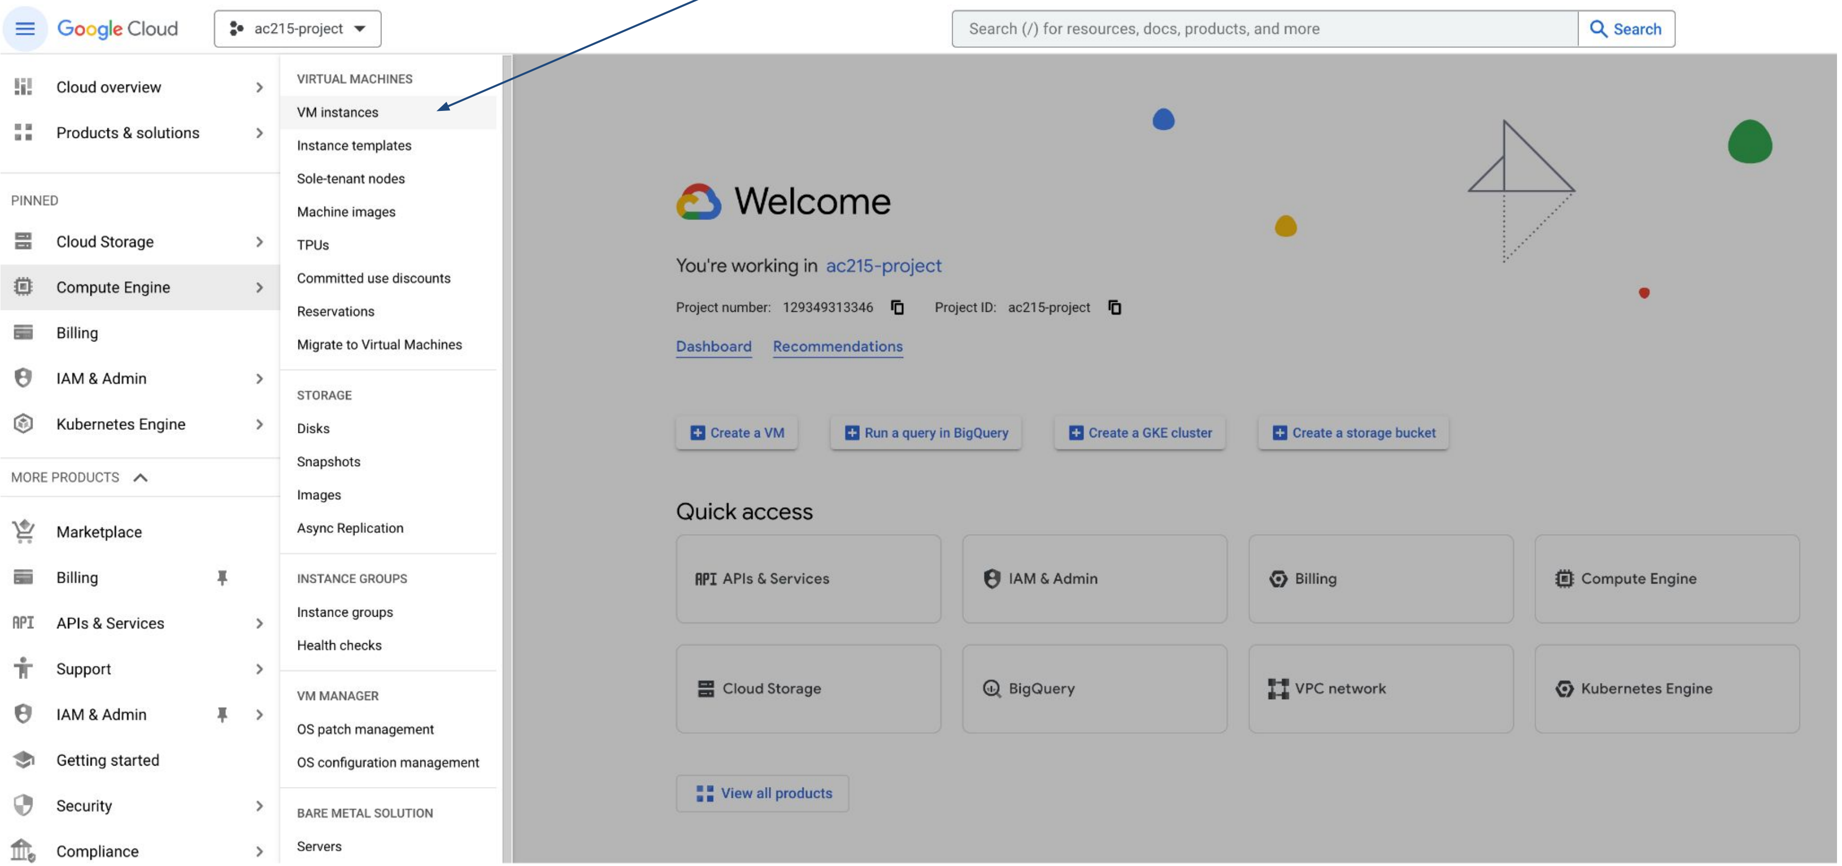

Virtual Machines Tutorial

Select Virtual Machines

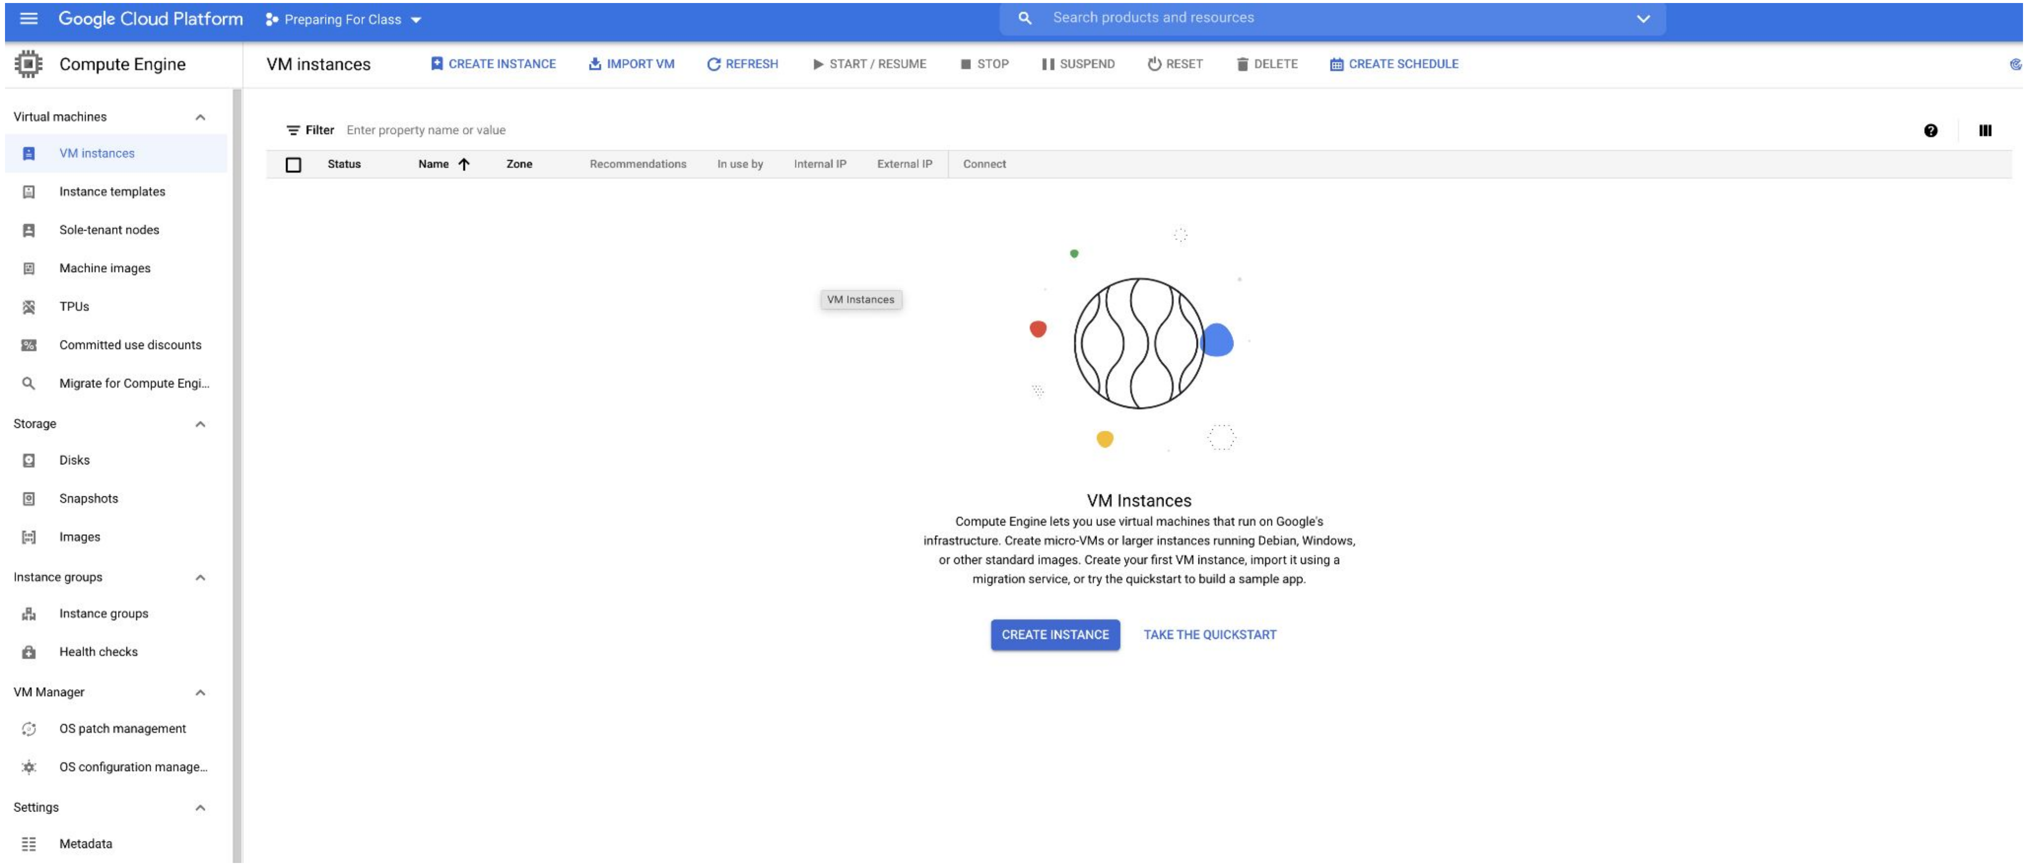

Virtual Machines Tutorial

Select all defaults

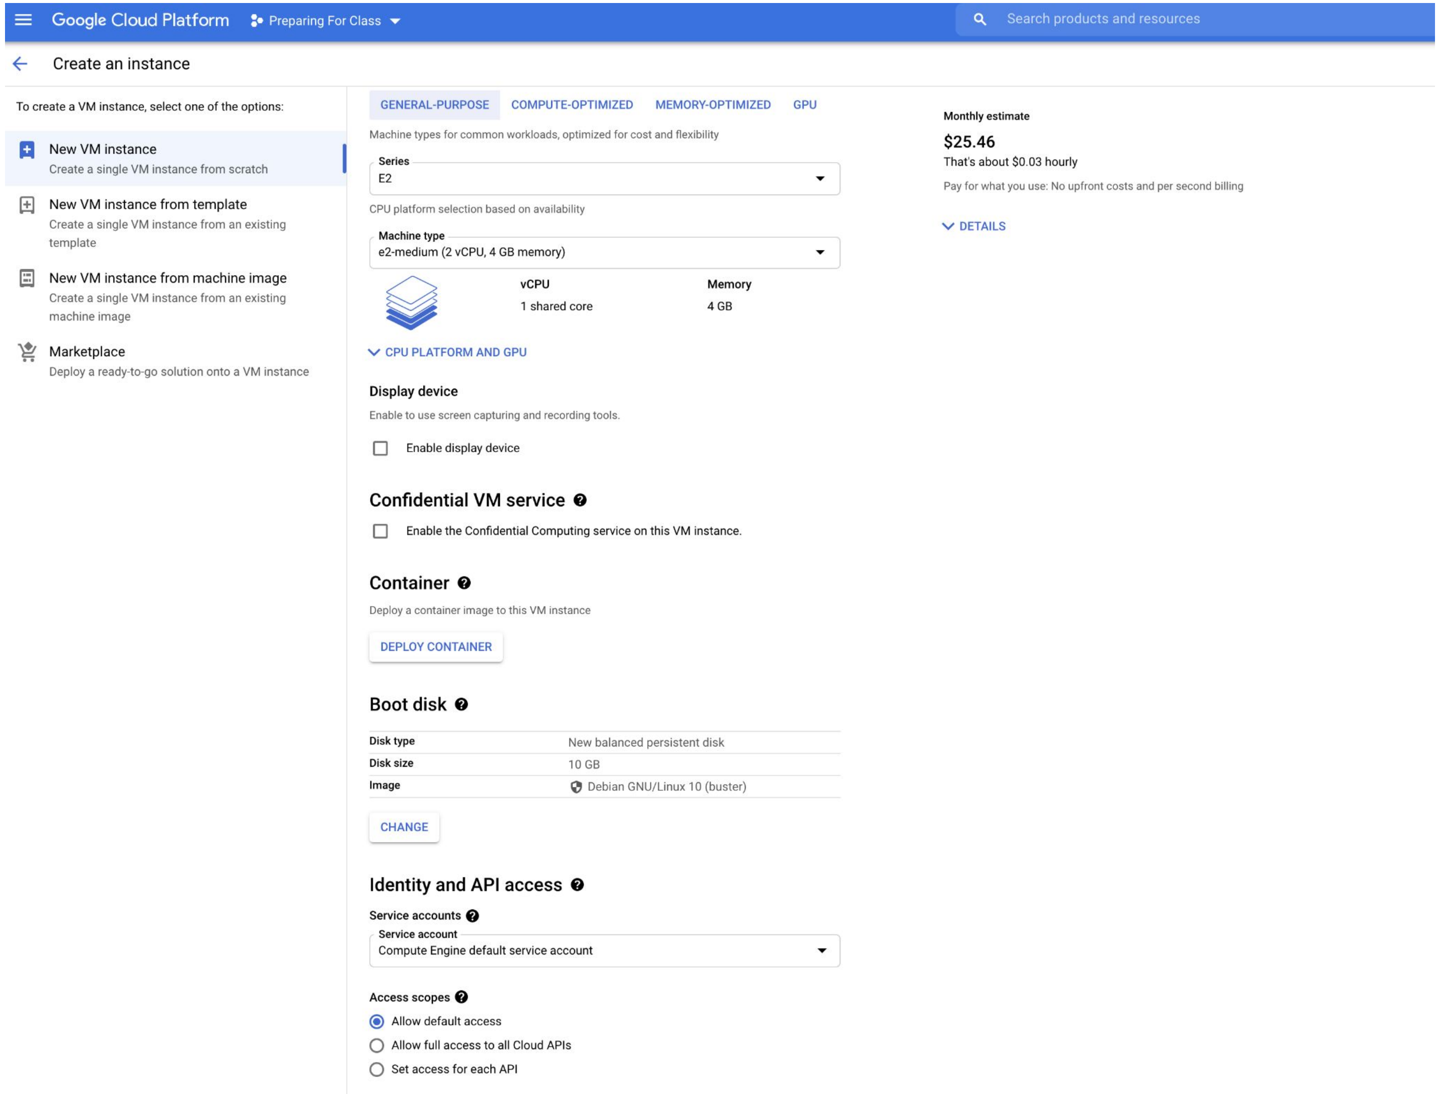

Virtual Machines Tutorial

Wait for instance to start and click on ssh

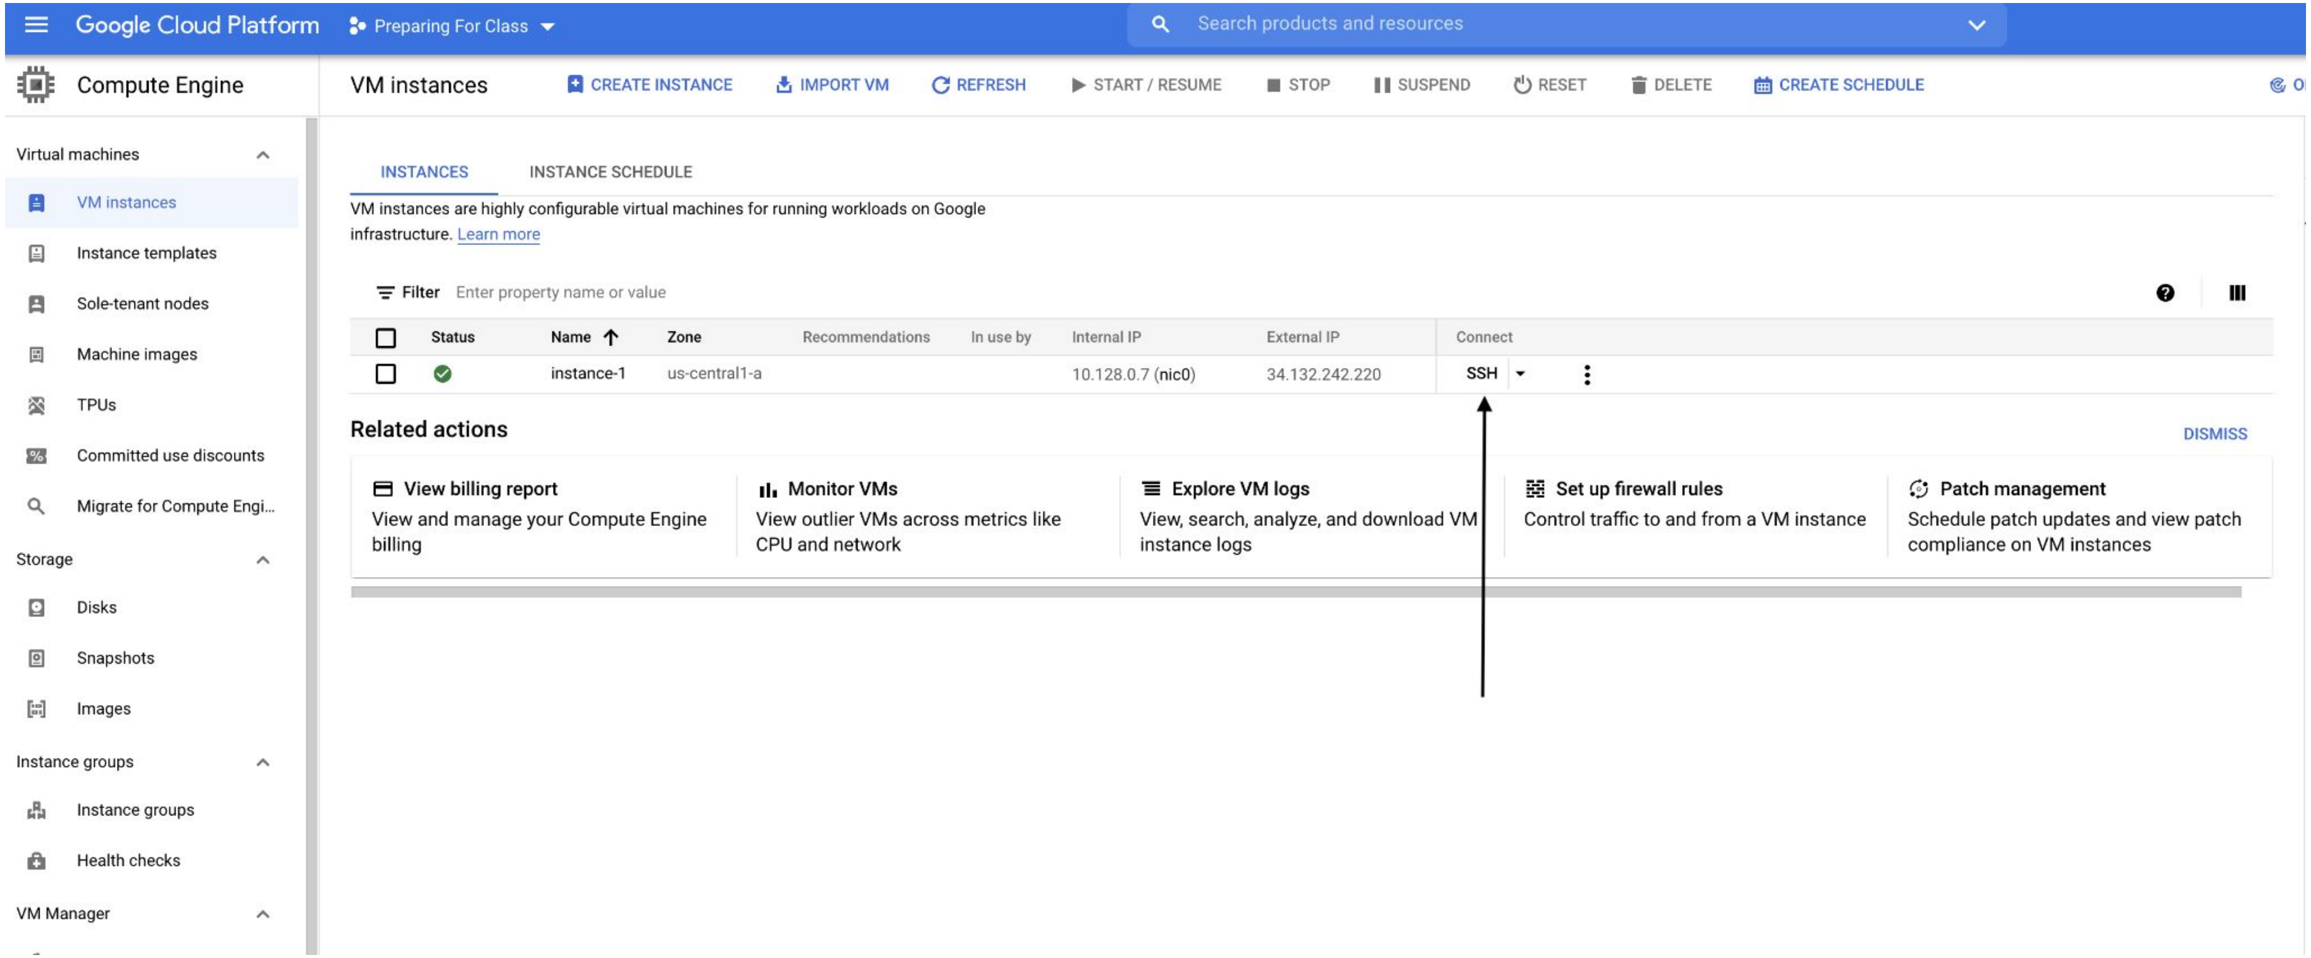

Virtual Machines Tutorial

And here is your virtual machine

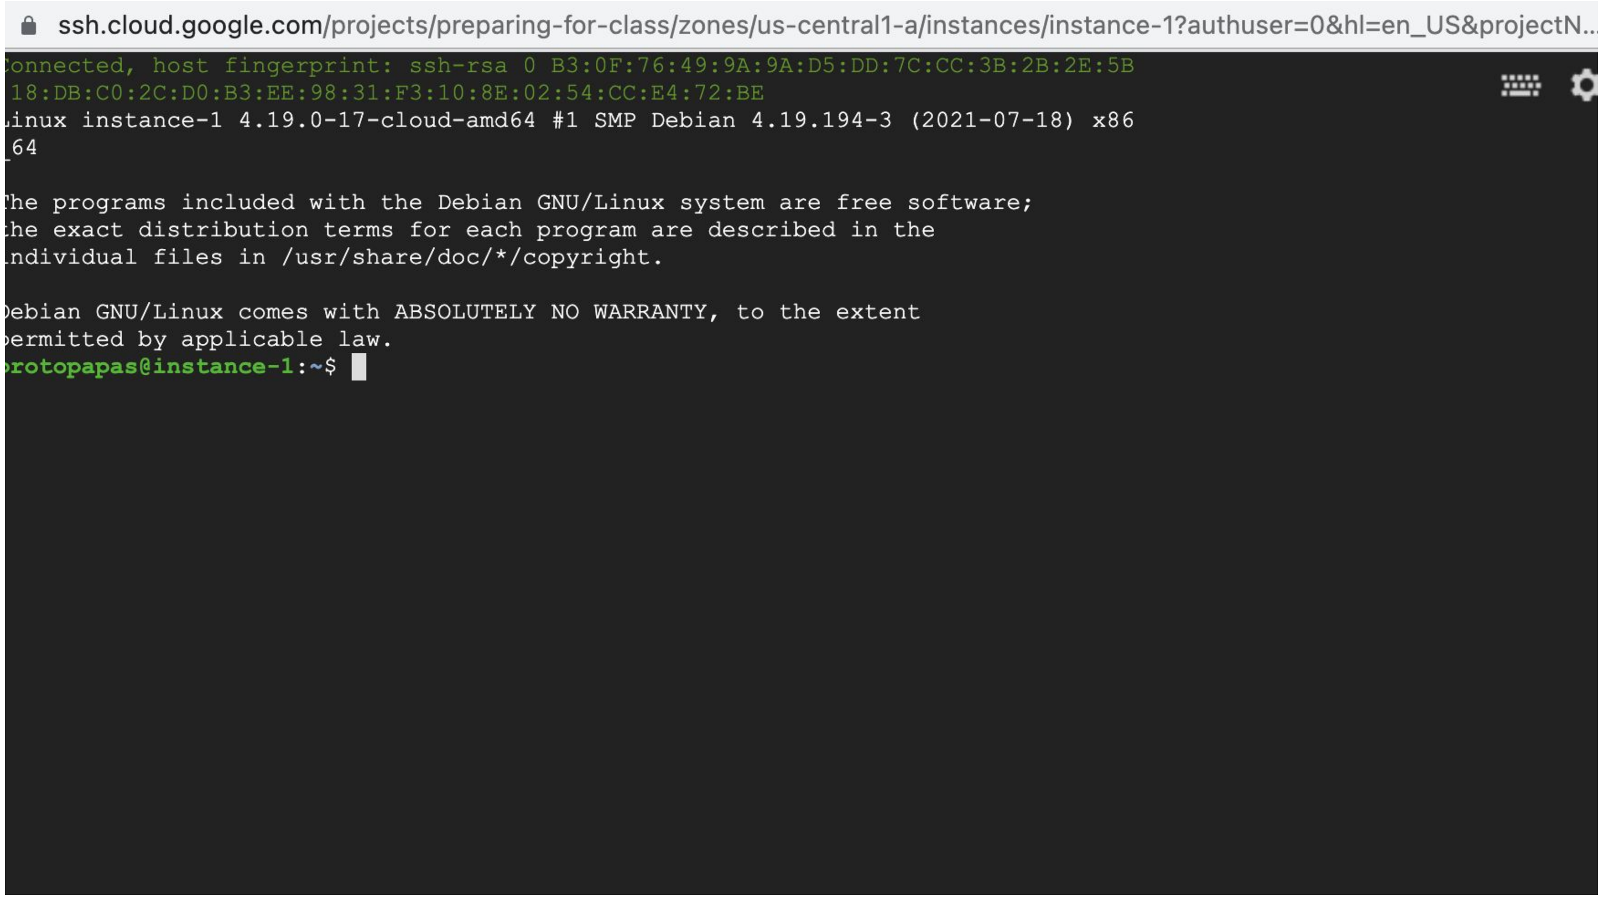

git clone https://github.com/dlops-io/simple-translate.git

Why should we use virtual machines?

Motivation

- **Uniform Operating Environments**: Desire for a standardized OS across all team member workstations.
- **Consistent Software Configuration:** Requirement for identical software setups across the team.
- **Effortless Instance Management:** The need for simple procedures to instantiate and terminate VMs.

Virtual Machines!

Key Components of Virtual Machines & Hypervisors?

- Virtual machines mimic real hardware like CPUs and hard drives.
- Hypervisors manage these virtual machines on a server.
- Unlimited VMs can be run, subject to hardware limits.
- The main OS is the "host," and VMs run "guest" OS.
- Guest VMs can have different operating systems.

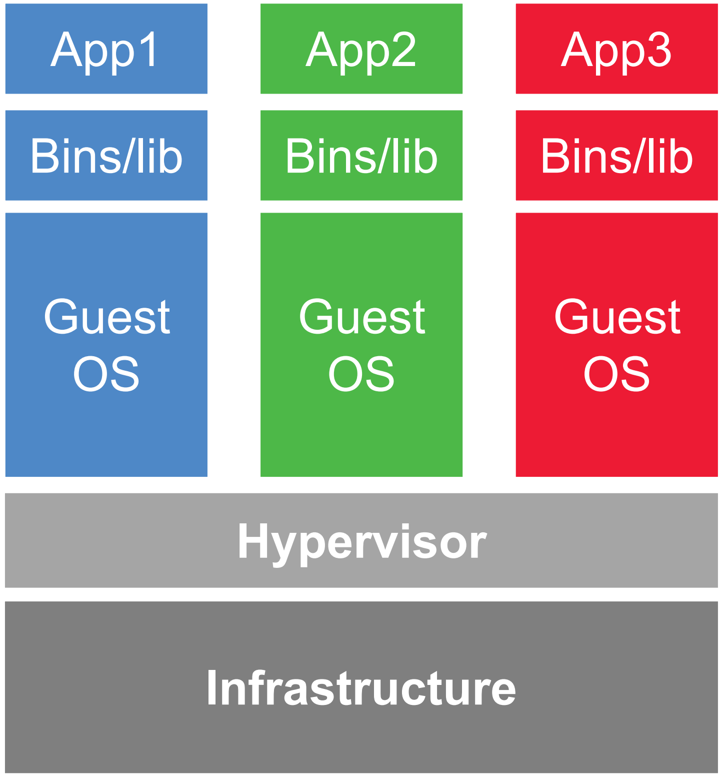

Machine Virtualization

Why should we use virtual machines?

Advantages

- **Complete Autonomy:** it works like a separate computer system; it is like running a computer within a computer.
- **Enhance Security**: the software inside the virtual machine cannot affect the actual computer.
- **Cost-Effectiveness:** Purchase a single machine and run multiple operating systems.
- **Widely Adopted:** Utilized by all major cloud providers for on-demand server instances.

Software for Virtualization

- VirtualBox
- VMWare
- Parallels

Why should we use virtual machines?

Limitations

- **Local Hardware Dependency:** Relies on the hardware resources of the host machine.
- **Limited Portability:** Large file sizes can impede easy transfer or deployment.
- **Resource Overhead:** Additional computational and memory resources are required to operate.
- **Reduced Performance:** The guest system typically runs slower than the host environment.
- **Slow Initialization:** Extended startup times compared to native systems.
- **Graphics Constraints:** May lack the graphical capabilities of the host system.

Outline

1. Motivation
2. Virtual Machines
3. **Virtual Environments**

What are virtual environments

A virtual environment is an isolated Python setting in which the interpreter can execute libraries and scripts independently of other virtual environments.

- Consider a virtual environment as a directory containing the following components:
  - `site_packages/`: *A directory where third-party libraries are installed.*
  - `Symlinks`: *Links to system executables.*
  - `Scripts`: *These ensure that the code utilizes the interpreter and site packages specific to the virtual environment.*

Why should we use virtual environment?

Maggie took CS109B and used to run her Jupyter notebooks from the Anaconda prompt. Whenever she installed a module, it was placed in one of the following folders: `bin`, `lib`, `share`, or `include`. She could then import the module and used it without any issue.

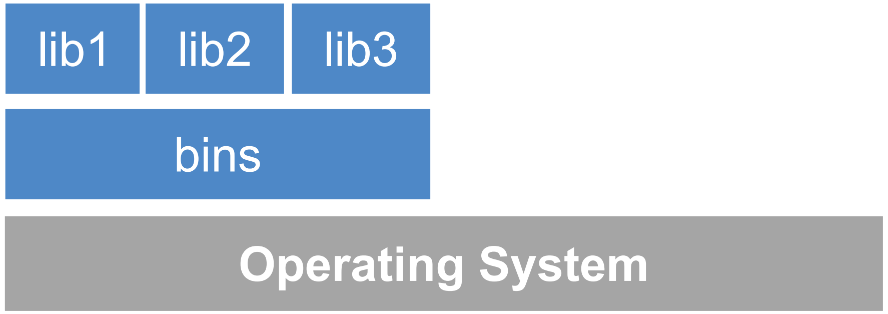

Maggie

```
$ which python
/c/Users/maggie/Anaconda3/python
```

Why should we use virtual environment?

Maggie begins taking AC215 and decides that isolating the new coding environment from previous ones would be beneficial to avoid package conflicts. To achieve this, she employs a layer of abstraction known as a virtual environment. This helps her keep modules organized and prevents issues while developing new projects.

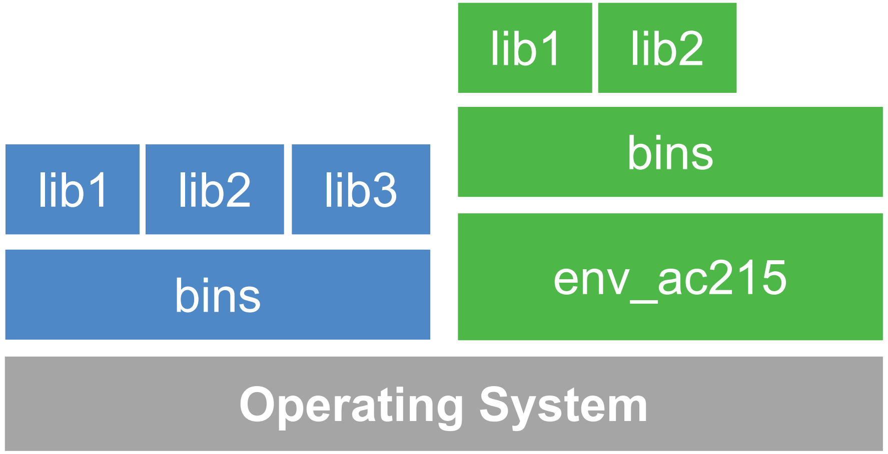

Maggie

```
$ which python
/c/Users/maggie/Anaconda3/envs/env_ac215/python
```

Why should we use virtual environment?

For the final project, Maggie collaborates with John and shares her working environment by distributing a .yml file for the Conda environment.

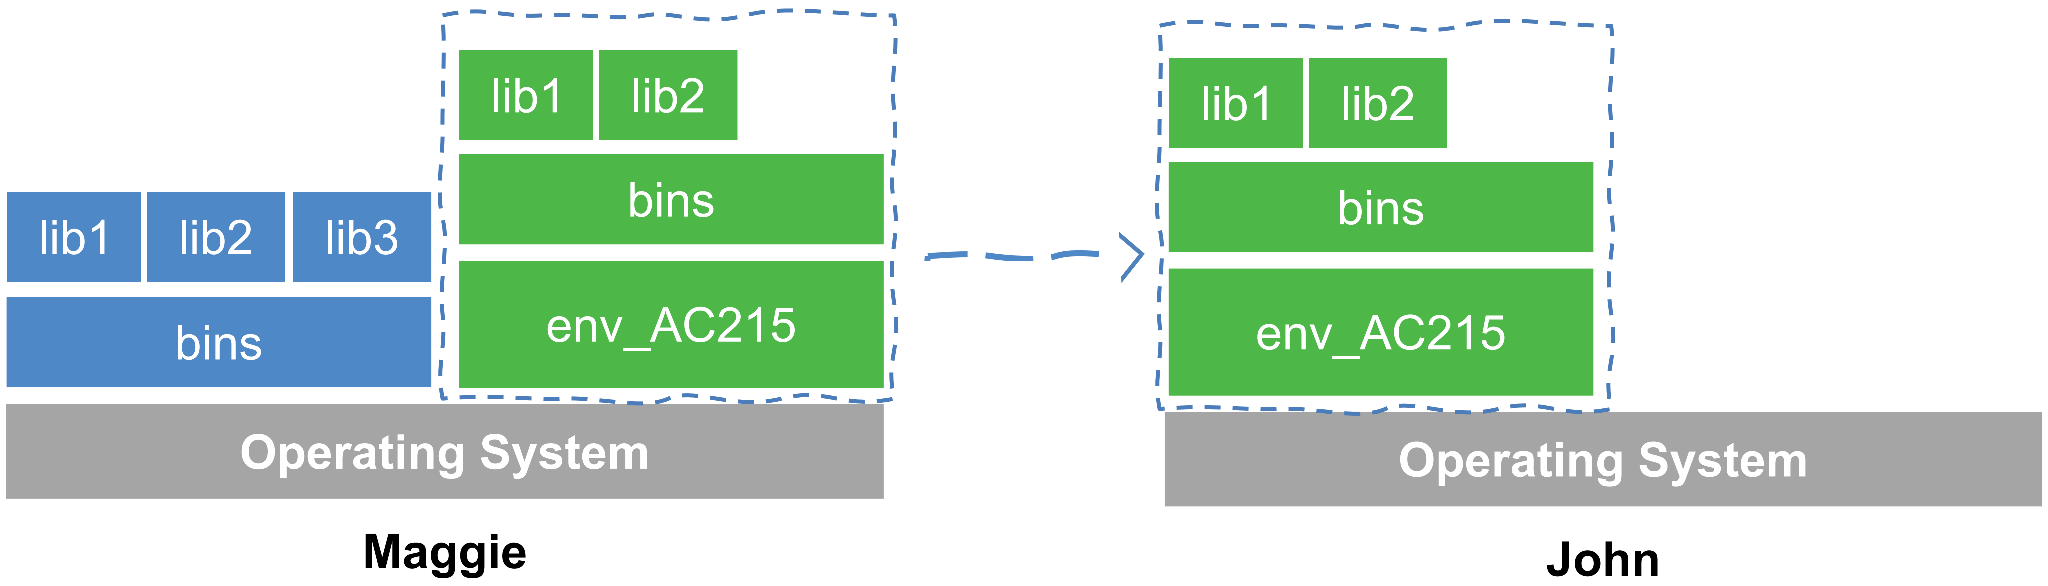

Why should we use virtual environment?

John experiments with a new method he learned in another class and adds a new library to the working environment. After seeing tremendous improvements, he sends Maggie back his code and a new .yml file (for conda env). She can now update her environment and replicate the experiment.

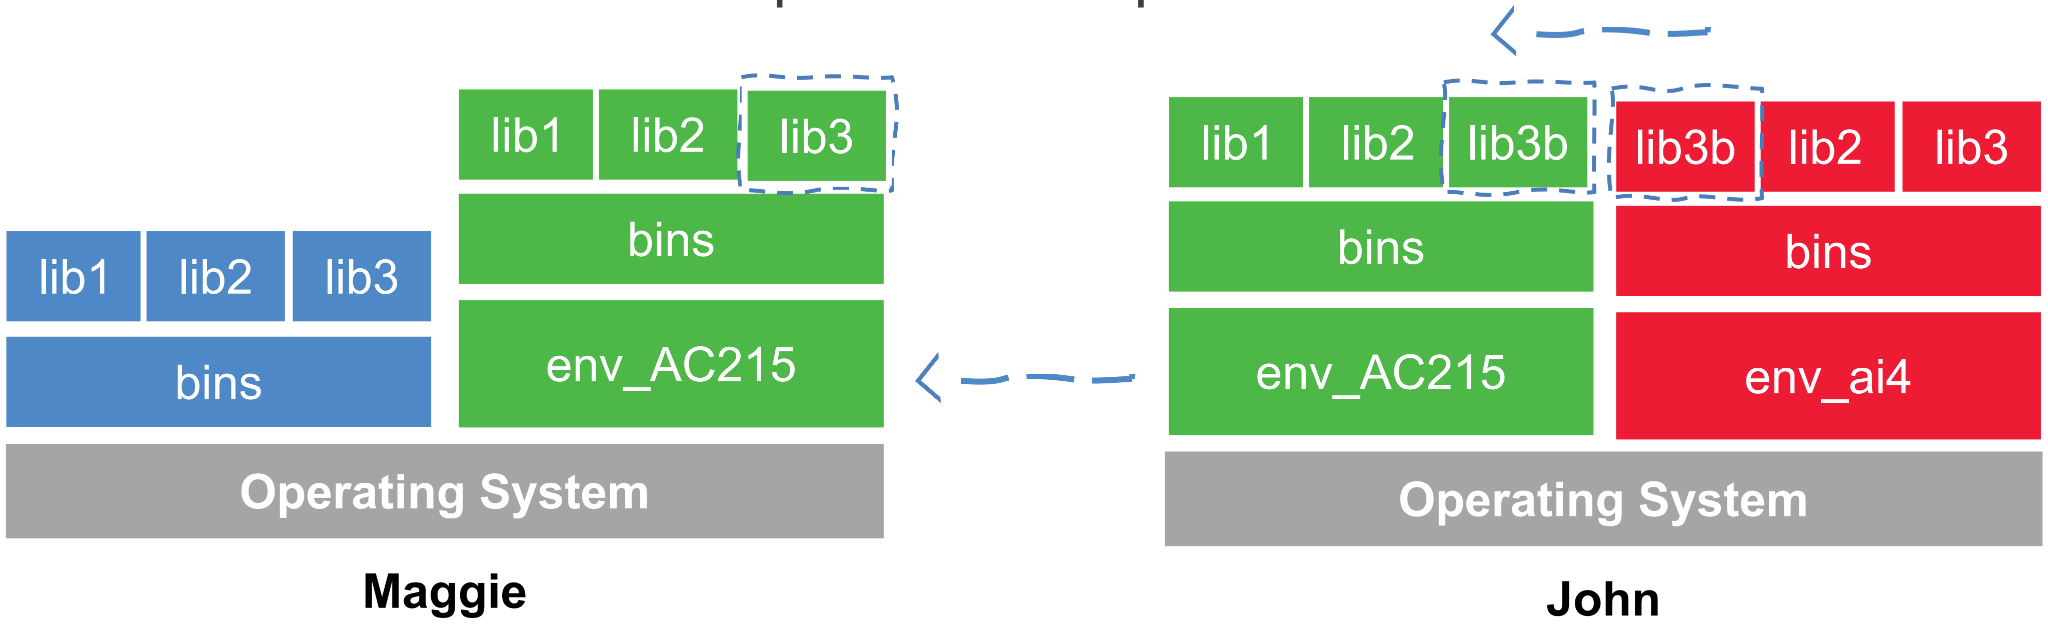

Why should we use virtual environment?

- What could go wrong?

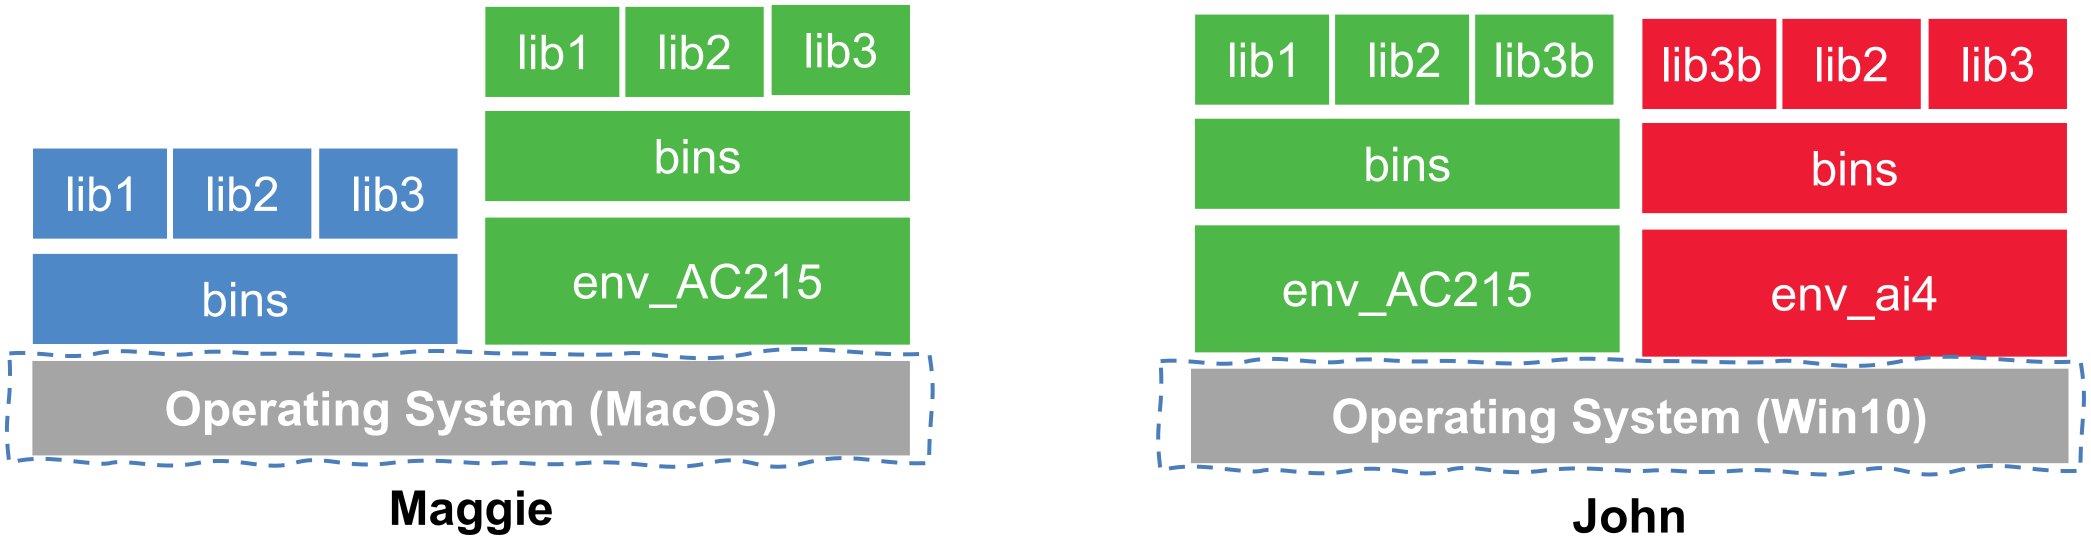

Why should we use virtual environment?

- What could go wrong?
- Unfortunately, Maggie and John are getting different results, which they suspect is due to their differing operating systems. Specifically, Maggie is using macOS, while John is on Windows 10.

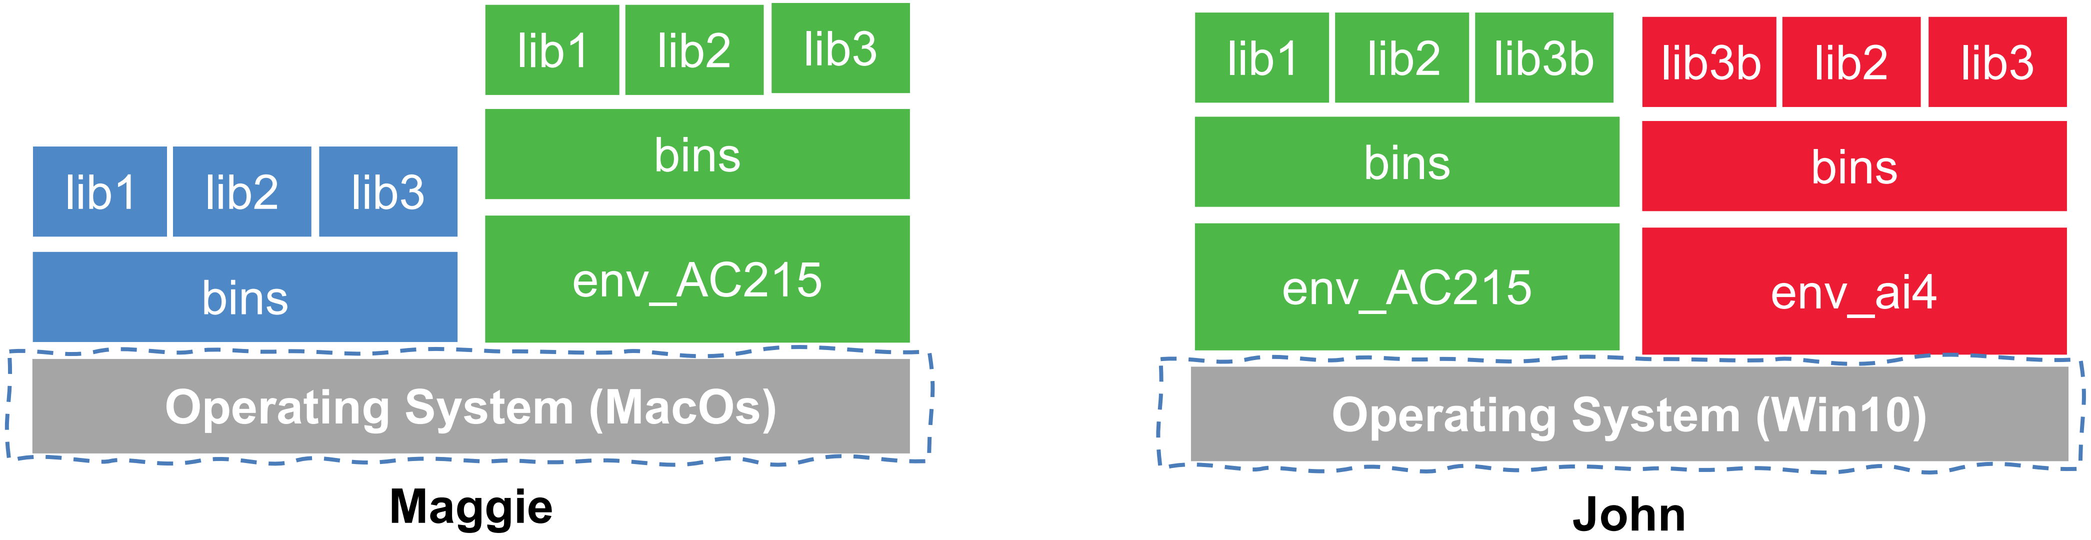

Why should we use virtual environment?

- Streamlines code development and usage.
- Isolates dependencies in separate "sandboxes" for easy switching between applications.
- Given an operating system and hardware, we can get the exact code environment set up using different technologies.

Virtual environments

Pros

- Reproducible Research: Enables consistent and replicable outcomes.
- Explicit Dependencies: Clearly defines all required software and packages.
- Enhanced Engineering Collaboration: Streamlines teamwork by standardizing environments.

Cons

- Setup Challenges: Initial environment configuration can be complex.
- Lack of Isolation: Does not completely isolate the working environment.
- OS Compatibility Issues: May not function consistently across different operating systems.

Creating Virtual Environments

- **virtualenv (python2) / venv (python3)**

  The default way to create virtual environments in python
- **conda**

  Is a package manager and environment manager for Data Scientists
- **pipenv**

  Production-ready tool that aims to bring the best of all packaging worlds to the Python world
- **mamba**

  Fast (C++) replacement for the Conda package manager that aims to offer quicker dependency resolution and installation - must do HW0 of CS109A

venv

- Virtual environments manager embedded in Python
- Incorporated into broader tools such as pipenv
- Allow to install modules using pip package manager

venv

How to use it:

- create an environment within your project folder `python3 -m venv your_env_name`
- it will add a folder called environment_name in your project directory
- activate environment: `source your_env_name/bin/activate`
- install requirements using: `pip install package_name=version`
- deactivate environment once done: `deactivate`

Conda

- Virtual environments manager embedded in Anaconda
- Allow to use both conda and pip to manage and install packages
- Base virtual environment comes pre-installed with various engineering and data science packages

Conda

How to use it:

- create an environment

  `conda create --name your_env_name python=3.7`

- it will add a folder located within your anaconda installation

  `/Users/your_username /anaconda3/envs/your_env_name`

- activate environment `conda activate your_env_name` (should appear in your shell)
- install requirements using `conda install package_name=version`
- deactivate environment once done `conda deactivate`
- duplicate your environment using YAML file `conda env export > my_environment.yml`
- to recreate the environment now use `conda env create -f environment.yml`

Conda

How to use it:

- find which environment you are using

  `conda env list`
- create an environment

  `conda create --name your_env_name python=3.7`
- it will add a folder located within your anaconda installation

  `/Users/your_username/[opt]/anaconda3/envs/your_env_name`
- activate environment

  `conda activate your_env_name` (should appear in your shell)
- install requirements using

  `conda install package_name=version`
- deactivate environment once done

  `conda deactivate`
- duplicate your environment using YAML file `conda env export > my_environment.yml`
- to recreate the environment now use `conda env create -f environment.yml`

PipEnv

- Built on top of *VirtualEnv*
- Fixes many shortcomings of *VirtualEnv*
- Distinguish development vs. production environments
- Automatically keeps track of packages and package dependencies using a `Pipfile` & `Pipfile.lock`

PipEnv

How to use it:

- Need to `pip install pipenv`
- To create a new environment run `pipenv install`
- Activate the environment by `pipenv shell`
- To install a new package `pipenv install numpy` or `pip install numpy` (this will not lock the package automatically)
- To sync from an existing Pipfile: `pipenv sync`

More on Virtual environments

Further readings

- Pipenv: Python Dev Workflow for Humans

  https://pipenv.pypa.io/en/latest/
- For detailed discussions on similarities and differences among virtualenv and conda

  https://jakevdp.github.io/blog/2016/08/25/conda-myths-and-misconceptions/
- More on venv and conda environments

  https://towardsdatascience.com/virtual-environments-104c62d48c54

  https://towardsdatascience.com/getting-started-with-python-environments-using-conda-32e9f2779307

Virtual Environments vs Virtual Machine

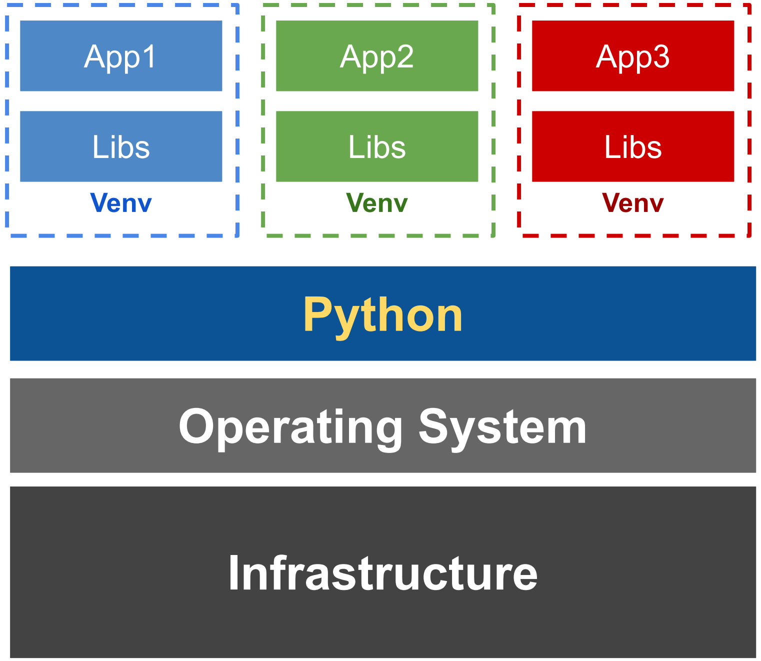

Virtual Environments

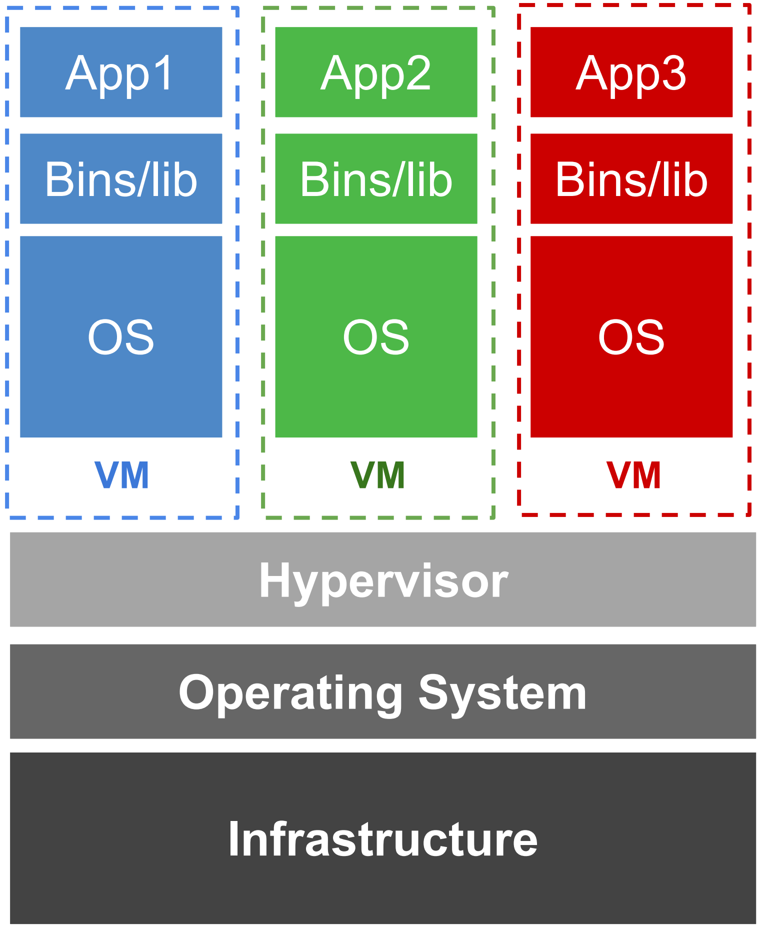

Virtualization

Virtual Environment Tutorial

- Let us run the simple-translate app using Virtual Environment
- For this we will do the following:
  - Create a VM Instance
  - SSH into the VM
  - Install dependencies: git, python
  - Download and run the simple-translate python script
- Full instructions can be found here

THANK YOU