Fastai Setup

jueves, 16 de febrero de 2023 08:14 a. m.

I followed this guide here https://forums.fast.ai/t/platform-local-server-ubuntu/65851 , I updated the versions but didn't run some parts as it seemed unneccessary.

wget https://repo.anaconda.com/archive/Anaconda3-2022.10-Linux-x86_64.sh

Extract and install the download using

bash Anaconda3-2022.10-Linux-x86_64.sh

You will have to press ENTER and space a few times to accept the licensing agreement. The install will need a new shell access, so exit out and get a new terminal (or reboot).

Execute the following commands to create the fastai environment.

```
conda create --name fastai
conda activate fastai
```

Run these commands to install fastai and any dependencies. Click y or yes or whatever as needed to add the items (from fastai github readme)

```
conda install -c pytorch -c fastai fastai
conda install pytorch==1.12.1 torchvision==0.13.1 torchaudio==0.12.1 cudatoolkit=11.6 -c pytorch -c conda-forge
```

NOTE: This pytorch version is the one that makes the gpu work since my CUDA version is 11.6. Without more recent versions, it apparently uses the processor.

```
conda uninstall --force jpeg libtiff -y
conda install -c conda-forge libjpeg-turbo
C="cc -mavx2" pip install --no-cache-dir -U --force-reinstall --no-binary :all: --compile pillow-simd
conda install jupyter notebook
conda install -c conda-forge jupyter_contrib_nbextensions
```

Copy the fastbook repo and start the jupyter notebook server

```
git clone https://github.com/fastai/fastbook
cd fastbook
Jupyter notebook --no-browser --ip 0.0.0.0 --port 8888
```

You should be able to access the fastbook jupyter notebooks and be able to train the model with the use of the GPU

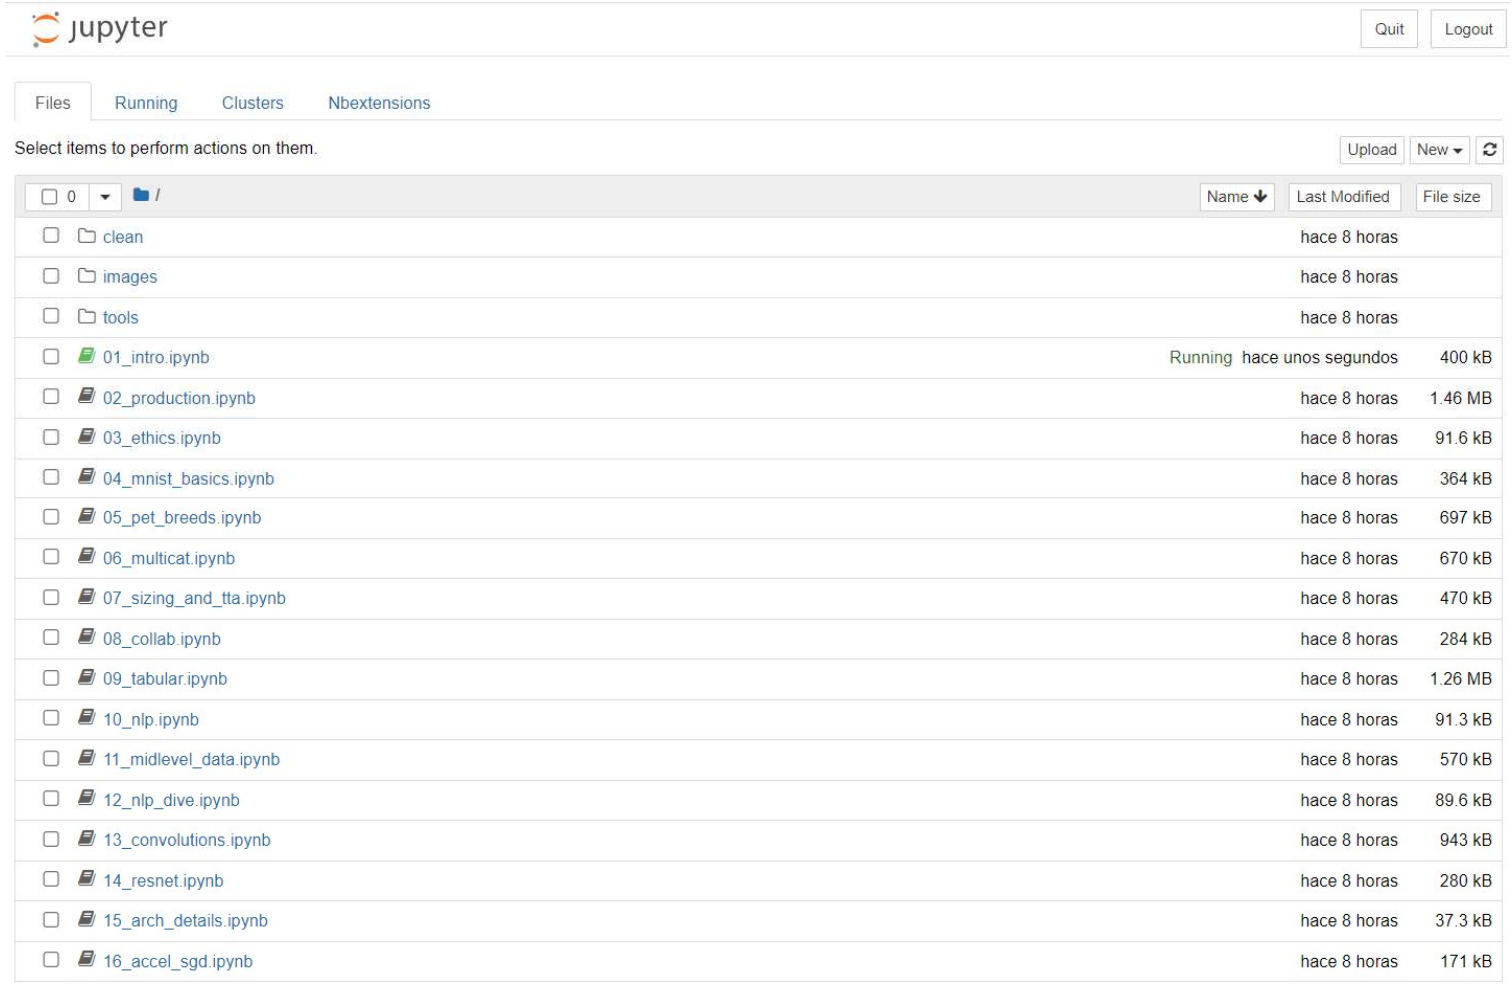

In [*]:

```
#id first_training
#caption Results from the first training
# CLICK ME
from fastai.vision.all import *
path = untar_data(URLs.PETS)/'images'

def is_cat(x): return x[0].isupper()
dls = ImageDataLoaders.from_name_func(
    path, get_image_files(path), valid_pct=0.2, seed=42,
    label_func=is_cat, item_tfms=Resize(224))

learn = vision_learner(dls, resnet34, metrics=error_rate)
learn.fine_tune(1)
```

```
/home/jgras/anaconda3/envs/fastai/lib/python3.10/site-packages/torchvision/models/_utils.py:208: UserWarning: The parameter
'pretrained' is deprecated since 0.13 and will be removed in 0.15, please use 'weights' instead.
  warnings.warn(
/home/jgras/anaconda3/envs/fastai/lib/python3.10/site-packages/torchvision/models/_utils.py:223: UserWarning: Arguments othe
r than a weight enum or `None` for 'weights' are deprecated since 0.13 and will be removed in 0.15. The current behavior is
equivalent to passing `weights=ResNet34_Weights.IMAGENET1K_V1`. You can also use `weights=ResNet34_Weights.DEFAULT` to get t
he most up-to-date weights.
  warnings.warn(msg)
```

0.00% [0/1 00:00<?]

epoch	train_loss	valid_loss	error_rate	time

32.61% [30/92 00:06<00:13 0.7174]

You can see the GPU was used by opening a terminal, running nvidia-smi and checking the Process ID

```
Every 1.0s: nvidia-smi

Thu Feb 16 14:36:49 2023
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 510.108.03   Driver Version: 510.108.03   CUDA Version: 11.6     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  NVIDIA GeForce ...  Off  | 00000000:01:00.0  On |                  N/A |
| 38%   47C    P2    97W / 170W |   4408MiB / 12288MiB |     94%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                                  |
|  GPU   GI   CI        PID   Type   Process name                  GPU Memory |
|        ID   ID                                                   Usage      |
|=============================================================================|
|    0   N/A  N/A    485991      C   ...a3/envs/fastai/bin/python     4405MiB |
+-----------------------------------------------------------------------------+
```

NOTE: The CUDA version is important, as I later found out that for CUDA Version 11.6 the pytorch version that would work is 1.12.1. If I used a more recent pytorch version, the code would still work, but it would be the processor doing the job, not the GPU.

Errors I ran into

When running the first jupyter notebook, the training went successfully. However, there was an error that appeared when an image was submitted, here's how I solved it.

```
learn = vision_learner(dls, resnet34, metrics=error_rate)
learn.fine_tune(1)
```

```
/home/jgras/anaconda3/envs/fastai/lib/python3.10/site-packages/torchvision/models/_utils.py:208: UserWarning: The parameter
'pretrained' is deprecated since 0.13 and will be removed in 0.15, please use 'weights' instead.
  warnings.warn(
/home/jgras/anaconda3/envs/fastai/lib/python3.10/site-packages/torchvision/models/_utils.py:223: UserWarning: Arguments othe
r than a weight enum or `None` for 'weights' are deprecated since 0.13 and will be removed in 0.15. The current behavior is
equivalent to passing `weights=ResNet34_Weights.IMAGENET1K_V1`. You can also use `weights=ResNet34_Weights.DEFAULT` to get t
he most up-to-date weights.
  warnings.warn(msg)
```

epoch	train_loss	valid_loss	error_rate	time
0	0.170623	0.035943	0.012855	00:17

epoch	train_loss	valid_loss	error_rate	time
0	0.048282	0.027205	0.008796	00:23

I then uploaded a .jpg image and verified it was uploaded correctly:

In [12]:
```
#hide_output
uploader = widgets.FileUpload()
uploader
#img = PILImage.create(uploader)
#img.to_thumb(192)
```

Upload (1)

Upload (0)

In [15]:
```
img = PILImage.create(uploader.data[0])
img.to_thumb(192)
```

Out[15]:

I then ran the model and it showed an EROR:

```
In [ ]:
#hide
# For the book, we can't actually click an upload button, so we fake it
uploader = SimpleNamespace(data = ['images/chapter1_cat_example.jpg'])
```

```
In [16]:
img = PILImage.create(uploader.data[0])
is_cat,_,probs = learn.predict(img)
print(f"Is this a cat?: {is_cat}.")
print(f"Probability it's a cat: {probs[1].item():.6f}")
```

```
File ~/anaconda3/envs/fastai/lib/python3.10/site-packages/PIL/Image.py:2962, in open(fp, mode, formats)
   2959     fp = io.BytesIO(fp.read())
   2960     exclusive_fp = True
-> 2962 prefix = fp.read(16)
   2964 preinit()
   2966 accept_warnings = []

File ~/anaconda3/envs/fastai/lib/python3.10/site-packages/PIL/Image.py:519, in Image.__getattr__(self, name)
    512     warnings.warn(
    513         "Image categories are deprecated and will be removed in Pillow 10 "
    514         "(2023-07-01). Use is_animated instead.",
    515         DeprecationWarning,
    516         stacklevel=2,
    517     )
    518     return self._category
--> 519 raise AttributeError(name)

AttributeError: read
```

NOTE: I skipped the following cell so the image I uploaded was the one used, but even if I run it and then the following code block it still fails

```
In [ ]:
#hide
# For the book, we can't actually click an upload button, so we fake it
uploader = SimpleNamespace(data = ['images/chapter1_cat_example.jpg'])
```

Still fails even if a fastbook image is used:

```
In [18]:
img = PILImage.create(uploader.data[0])
is_cat,_,probs = learn.predict(img)
print(f"Is this a cat?: {is_cat}.")
print(f"Probability it's a cat: {probs[1].item():.6f}")
```

```
File ~/anaconda3/envs/fastai/lib/python3.10/site-packages/PIL/Image.py:2962, in open(fp, mode, formats)
   2959     fp = io.BytesIO(fp.read())
   2960     exclusive_fp = True
-> 2962 prefix = fp.read(16)
   2964 preinit()
   2966 accept_warnings = []

File ~/anaconda3/envs/fastai/lib/python3.10/site-packages/PIL/Image.py:519, in Image.__getattr__(self, name)
    512     warnings.warn(
    513         "Image categories are deprecated and will be removed in Pillow 10 "
    514         "(2023-07-01). Use is_animated instead.",
    515         DeprecationWarning,
    516         stacklevel=2,
    517     )
    518     return self._category
--> 519 raise AttributeError(name)

AttributeError: read
```

I searched online and found thisin stackoverflow https://stackoverflow.com/questions/51296492/python-attributeerror-while-reading-images-with-pil-library
After reading it, the fix was to remove the PILImage.create() function and it worked:

In [8]:
```
img = PILImage.create(uploader.data[0])
img.to_thumb(192)
```

Out[8]:

Now you can pass the uploaded file to the model. Make sure that it is a clear photo of a single dog or a cat, and not a line drawing, cartoon, or similar. The notebook will tell you whether it thinks it is a dog or a cat, and how confident it is. Hopefully, you'll find that your model did a great job:

In [17]:
```
#hide
# For the book, we can't actually click an upload button, so we fake it
uploader = SimpleNamespace(data = ['images/chapter1_cat_example.jpg'])
```

In [9]:
```
img = uploader.data[0]
is_cat,_,probs = learn.predict(img)
print(f"Is this a cat?: {is_cat}.")
print(f"Probability it's a cat: {probs[1].item():.6f}")
```

```
Is this a cat?: False.
Probability it's a cat: 0.028015
```