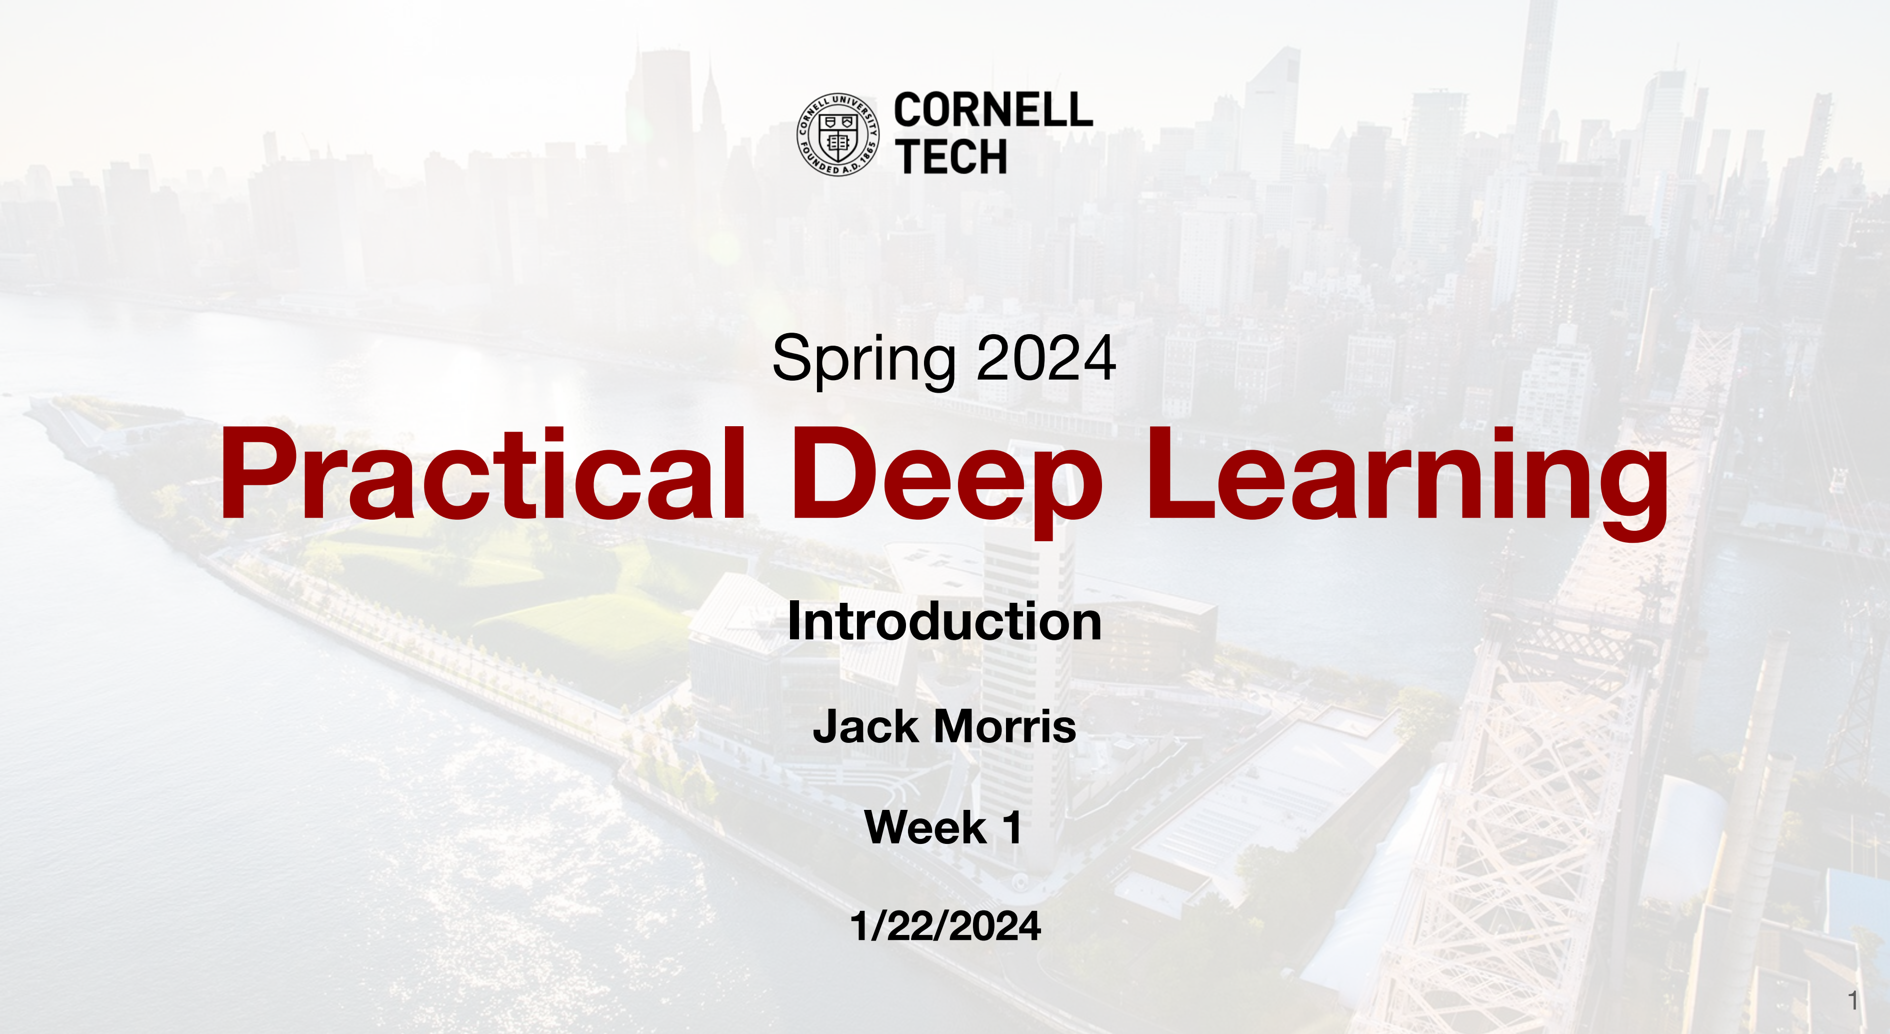

CORNELL
TECH

Spring 2024

Practical Deep Learning

Introduction

Jack Morris

Week 1

1/22/2024

About this class

- We'll meet every Monday for 8 weeks
- No assignments, just a project at the end
- No Zoom option – please pay attention :)
- Course website https://jxmo.io/deep-learning-workshop/
- Also please ask questions on Canvas

About the instructor

Jack Morris

jxm3@cornell.edu

PhD Student
Research Interest: NLP
(language models, text embeddings, inversion)

jxmo.io / Twitter @jxmnop

(Office Hours TBD)

Notes

- Anonymous feedback link: bit.ly/pdl24feedback
- Laptops are allowed (but please be respectful!)
- Will put my slides on course website

About me

Cornell
Ithaca, NY

Cornell Tech
New York, NY

Arlington, VA

University of Virginia
Charlottesville, VA

Google
Mountain View, CA

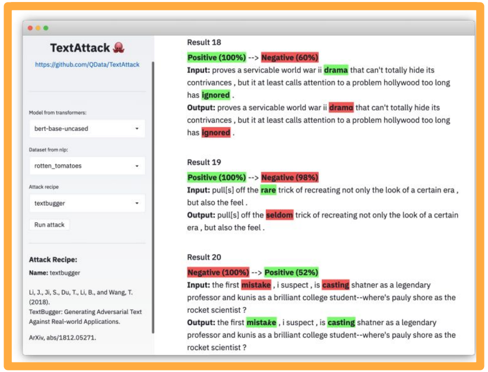

Model from transformers: bert-base-uncased

Dataset from nlp: rotten_tomatoes

Attack recipe: textbugger

Attack Recipe:

Name: textbugger

Li, J., Ji, S., Du, T., Li, B., and Wang, T. (2018).
TextBugger: Generating Adversarial Text Against Real-world Applications.

ArXiv, abs/1812.05271.

Result 18

Positive (100%) --> Negative (60%)

Input: proves a servicable world war ii **drama** that can't totally hide its contrivances , but it at least calls attention to a problem hollywood too long has **ignored** .

Output: proves a servicable world war ii **dramα** that can't totally hide its contrivances , but it at least calls attention to a problem hollywood too long has **ignored** .

Result 19

Positive (100%) --> Negative (98%)

Input: pull[s] off the **rare** trick of recreating not only the look of a certain era , but also the feel .

Output: pull[s] off the **seldom** trick of recreating not only the look of a certain era , but also the feel .

Result 20

Negative (100%) --> Positive (52%)

Input: the first **mistake** , i suspect , is **casting** shatner as a legendary professor and kunis as a brilliant college student--where's pauly shore as the rocket scientist ?

Output: the first **mistake** , i suspect , is **casting** shatner as a legendary professor and kunis as a brilliant college student--where's pauly shore as the rocket scientist ?

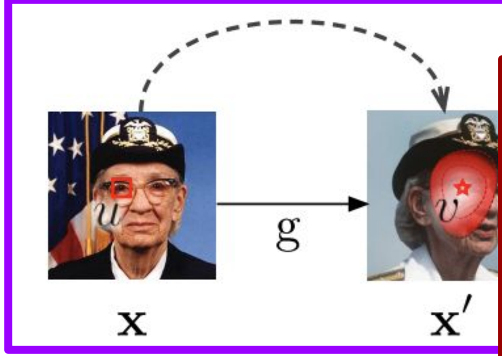

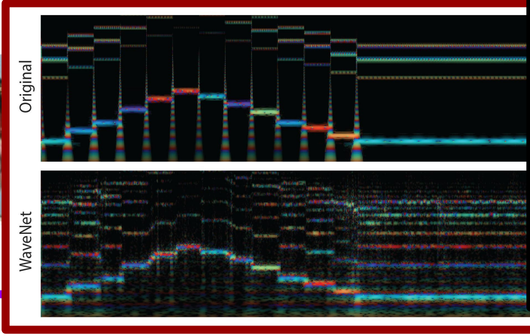

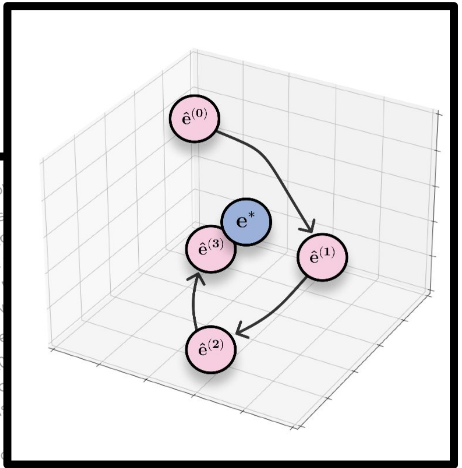

Semester overview

Deep learning toolstack ⚒️

Puzzle

Semester overview (and survey)

Deep learning toolstack ⚒️

Puzzle

Semester overview **(and survey)**

Deep learning toolstack ⚒️

Puzzle

Survey

bit.ly/pdl24survey

Schedule

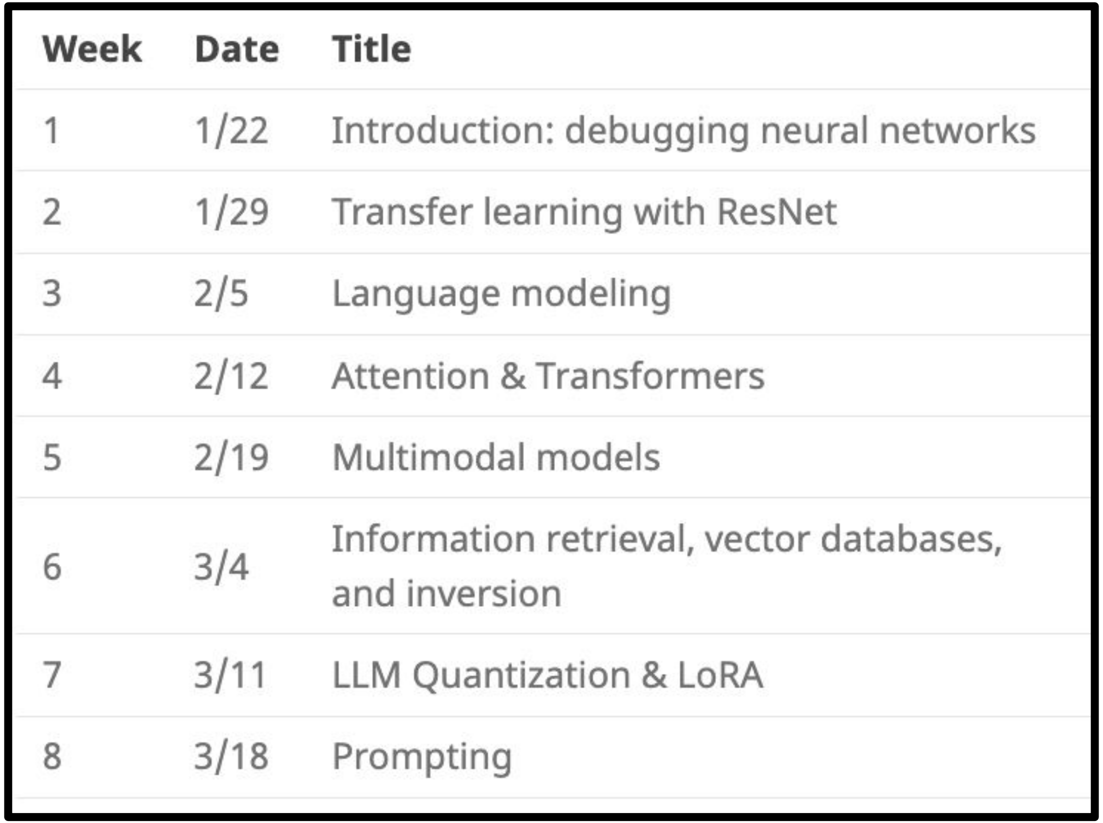

Week	Date	Title
1	1/22	Introduction: debugging neural networks
2	1/29	Transfer learning with ResNet
3	2/5	Language modeling
4	2/12	Attention & Transformers
5	2/19	Multimodal models
6	3/4	Information retrieval, vector databases, and inversion
7	3/11	LLM Quantization & LoRA
8	3/18	Prompting

(⅛) Debugging neural networks

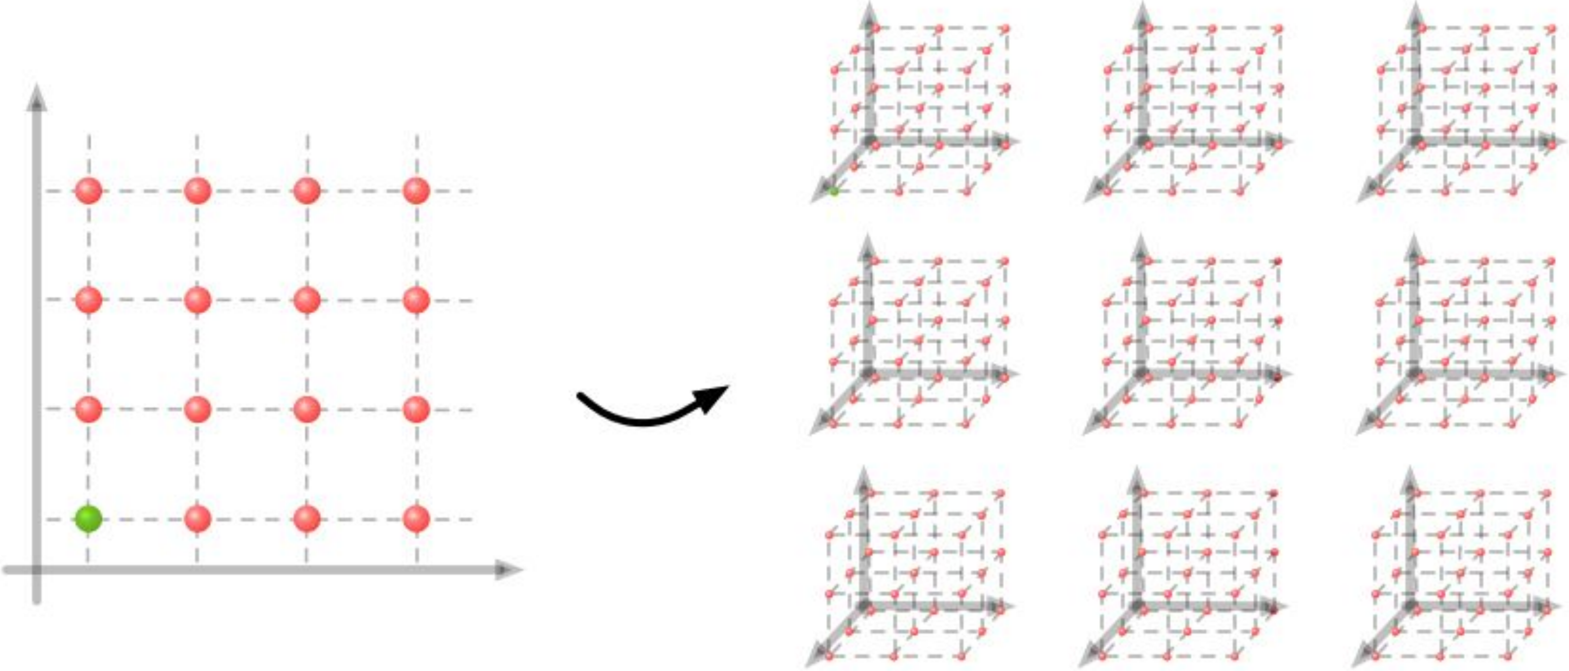

Software Development

Machine Learning

(2/8) Transfer learning

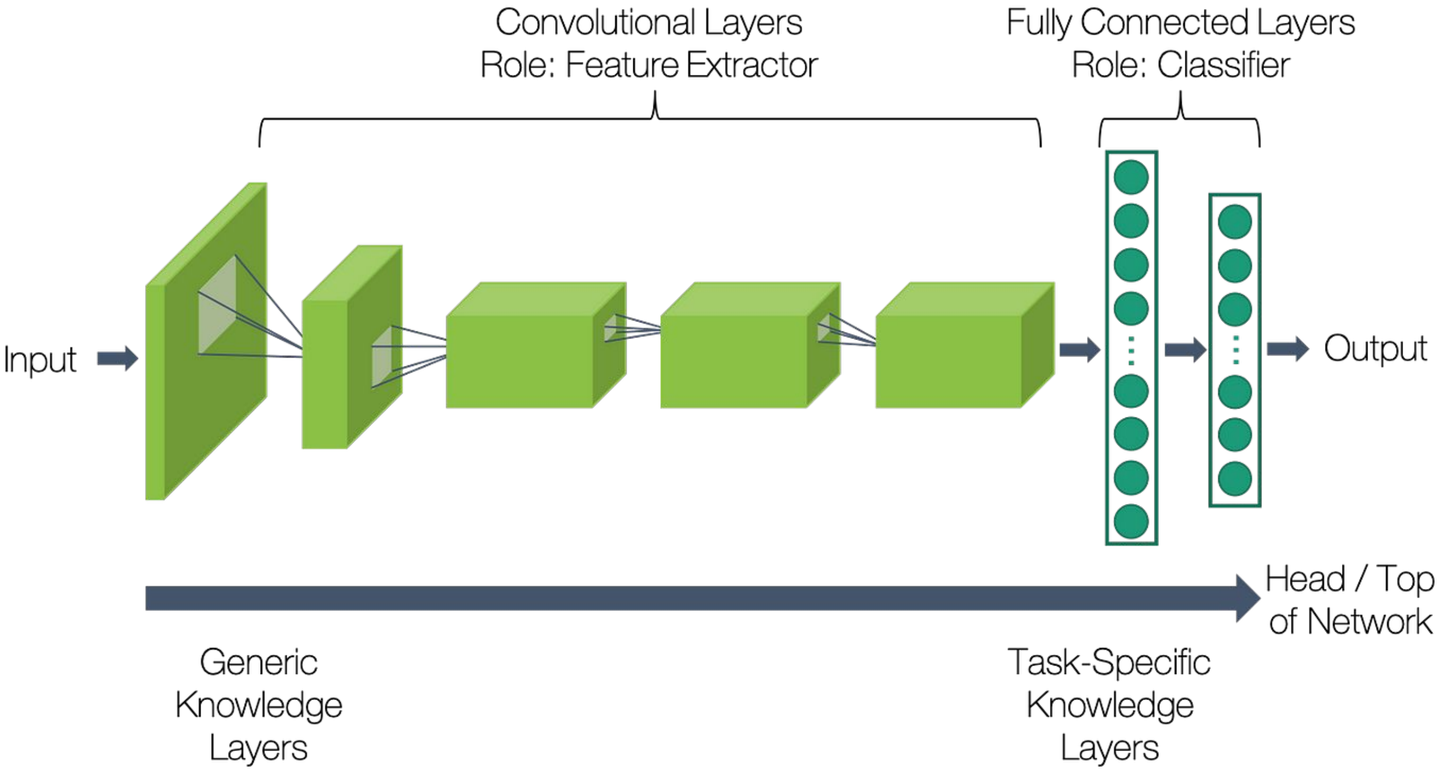

(⅜) Language modeling

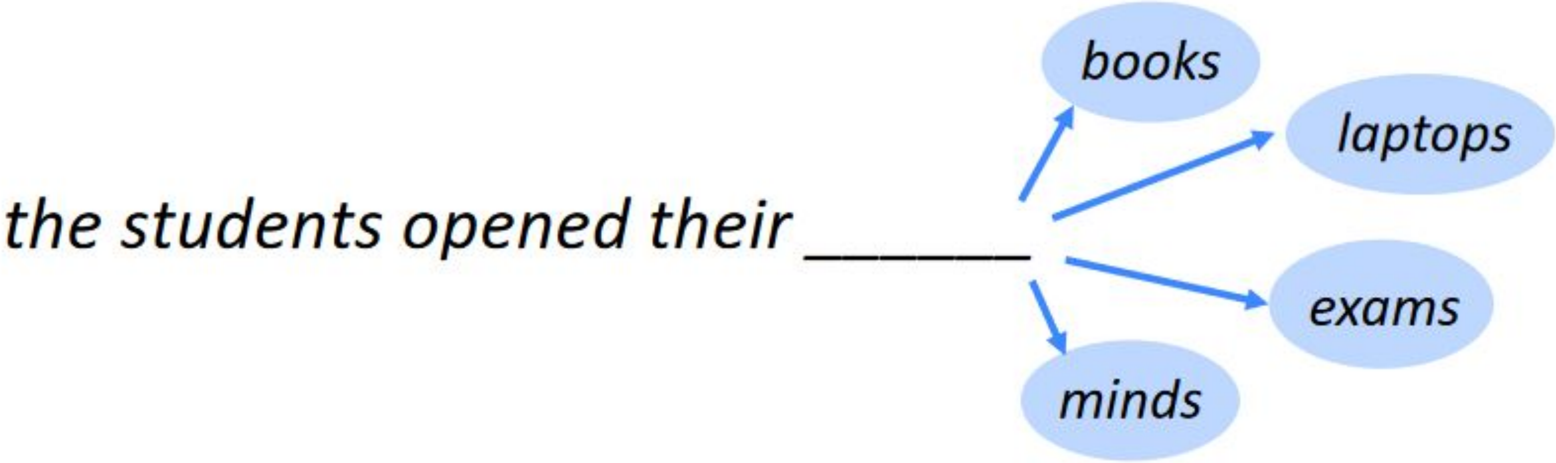

(4/8) Transformers

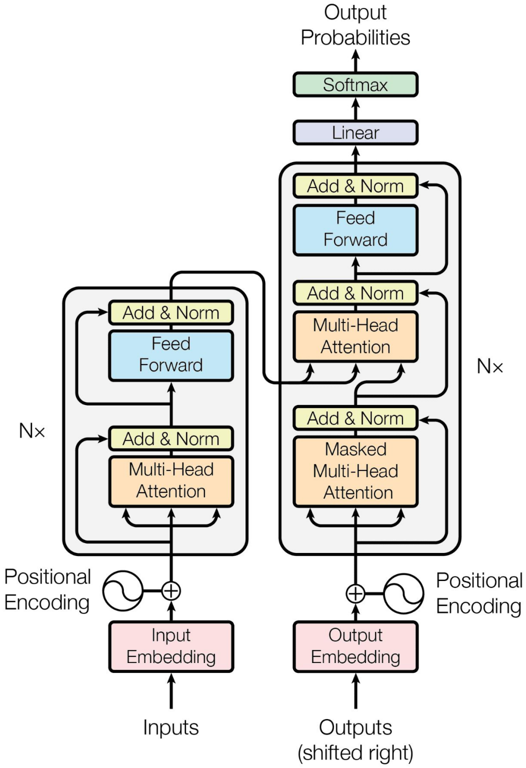

(5/8) Multimodal models

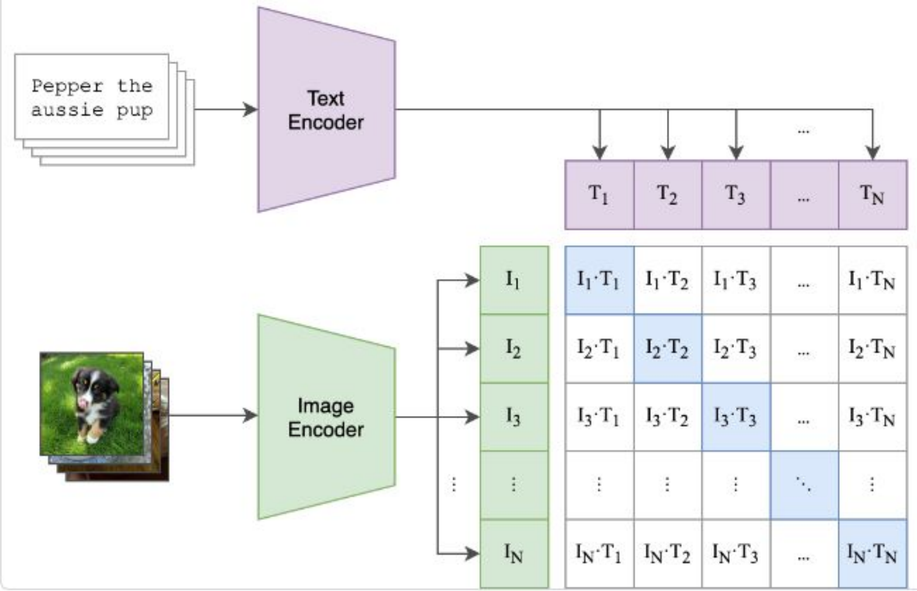

(6/8) Embeddings

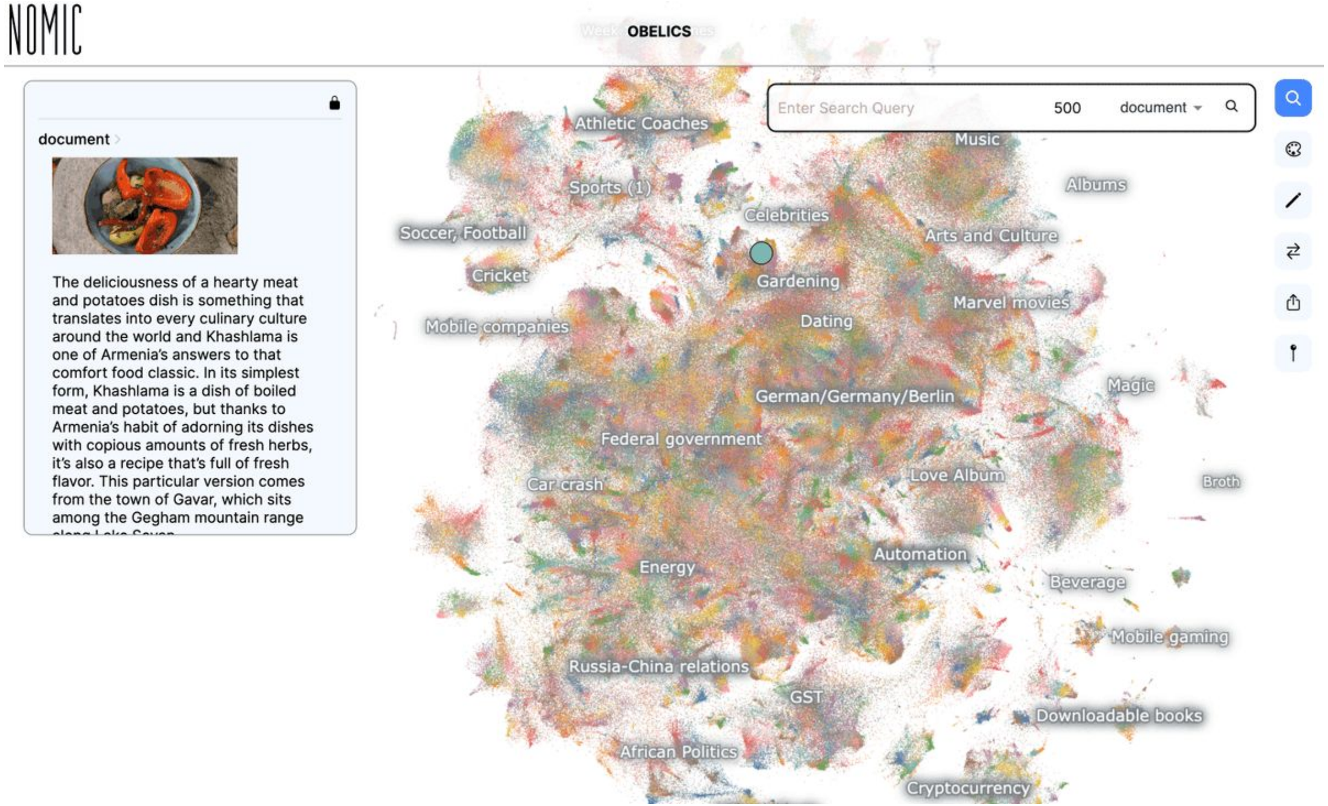

The deliciousness of a hearty meat and potatoes dish is something that translates into every culinary culture around the world and Khashlama is one of Armenia's answers to that comfort food classic. In its simplest form, Khashlama is a dish of boiled meat and potatoes, but thanks to Armenia's habit of adorning its dishes with copious amounts of fresh herbs, it's also a recipe that's full of fresh flavor. This particular version comes from the town of Gavar, which sits among the Gegham mountain range

(7/8) Local LLMs

(8/8) Prompting

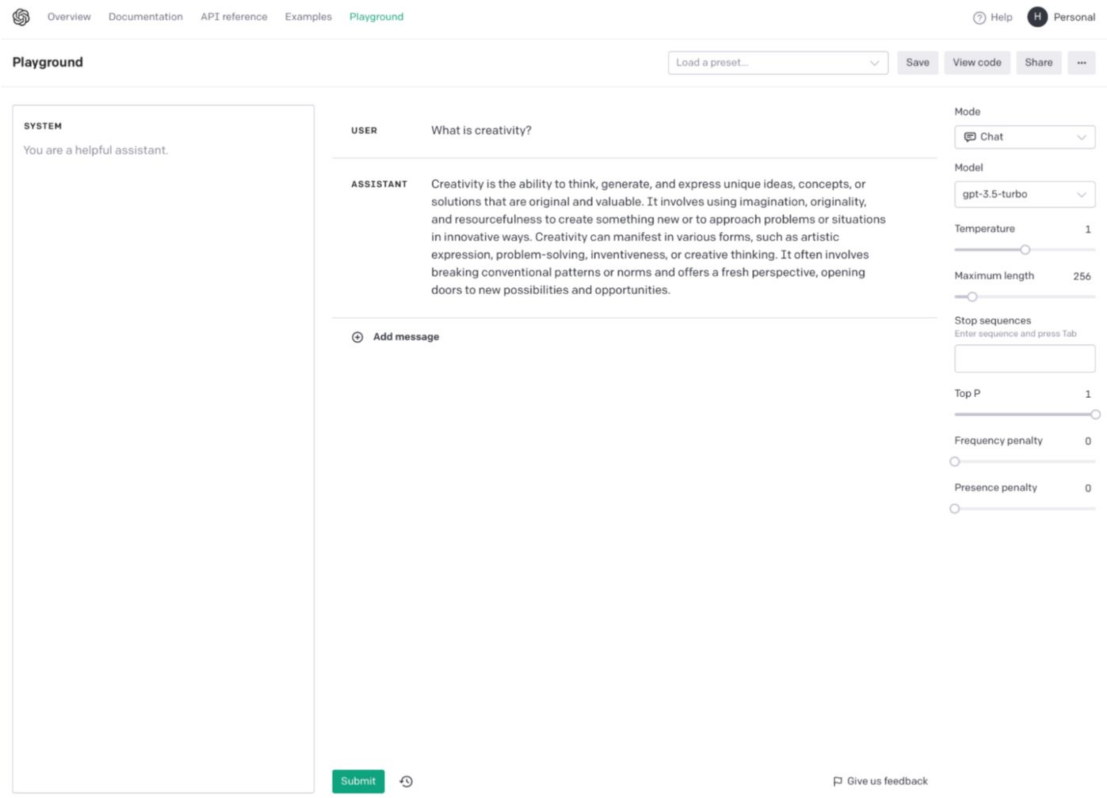

SYSTEM
You are a helpful assistant.

USER What is creativity?

ASSISTANT Creativity is the ability to think, generate, and express unique ideas, concepts, or solutions that are original and valuable. It involves using imagination, originality, and resourcefulness to create something new or to approach problems or situations in innovative ways. Creativity can manifest in various forms, such as artistic expression, problem-solving, inventiveness, or creative thinking. It often involves breaking conventional patterns or norms and offers a fresh perspective, opening doors to new possibilities and opportunities.

Mode: Chat
Model: gpt-3.5-turbo
Temperature: 1
Maximum length: 256
Stop sequences: Enter sequence and press Tab
Top P: 1
Frequency penalty: 0
Presence penalty: 0

Semester overview

Deep learning toolstack ⚒️

Puzzle

Deep learning toolstack ⚒️

What do we need to learn how to do?

- Read, write, and publish code
- Obtain and preprocess data
- Do computations required to run the model (lots of math)
- Run said computations on special hardware (GPUs)
- Understand results through visualization

Unix & Terminal

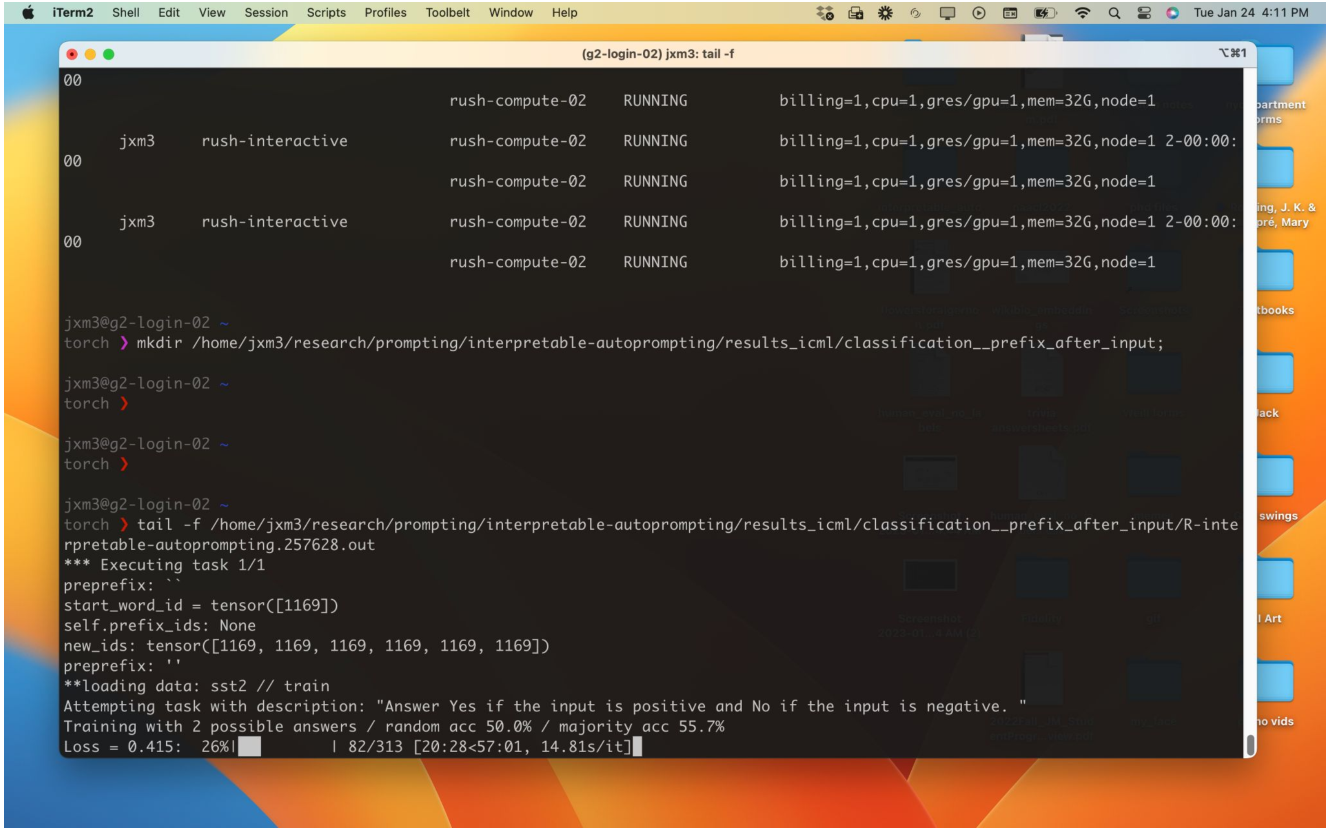

```
00
                                          rush-compute-02    RUNNING          billing=1,cpu=1,gres/gpu=1,mem=32G,node=1

      jxm3     rush-interactive           rush-compute-02    RUNNING          billing=1,cpu=1,gres/gpu=1,mem=32G,node=1 2-00:00:
00
                                          rush-compute-02    RUNNING          billing=1,cpu=1,gres/gpu=1,mem=32G,node=1

      jxm3     rush-interactive           rush-compute-02    RUNNING          billing=1,cpu=1,gres/gpu=1,mem=32G,node=1 2-00:00:
00
                                          rush-compute-02    RUNNING          billing=1,cpu=1,gres/gpu=1,mem=32G,node=1


jxm3@g2-login-02 ~
torch ❯ mkdir /home/jxm3/research/prompting/interpretable-autoprompting/results_icml/classification__prefix_after_input;

jxm3@g2-login-02 ~
torch ❯

jxm3@g2-login-02 ~
torch ❯

jxm3@g2-login-02 ~
torch ❯ tail -f /home/jxm3/research/prompting/interpretable-autoprompting/results_icml/classification__prefix_after_input/R-inte
rpretable-autoprompting.257628.out
*** Executing task 1/1
preprefix: ``
start_word_id = tensor([1169])
self.prefix_ids: None
new_ids: tensor([1169, 1169, 1169, 1169, 1169, 1169])
preprefix: ''
**loading data: sst2 // train
Attempting task with description: "Answer Yes if the input is positive and No if the input is negative. "
Training with 2 possible answers / random acc 50.0% / majority acc 55.7%
Loss = 0.415:  26%|██        | 82/313 [20:28<57:01, 14.81s/it]
```

VSCode

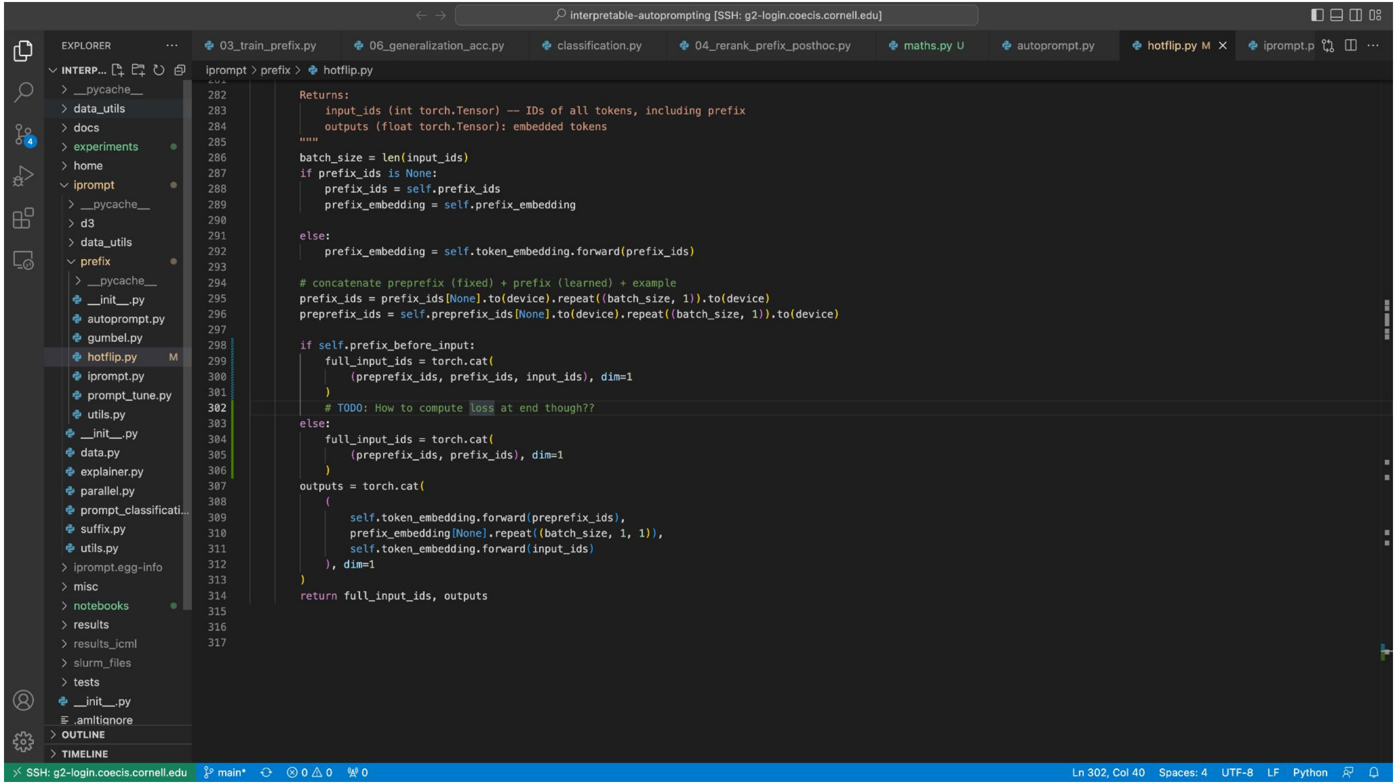

```python
        Returns:
            input_ids (int torch.Tensor) -- IDs of all tokens, including prefix
            outputs (float torch.Tensor): embedded tokens
        """
        batch_size = len(input_ids)
        if prefix_ids is None:
            prefix_ids = self.prefix_ids
            prefix_embedding = self.prefix_embedding

        else:
            prefix_embedding = self.token_embedding.forward(prefix_ids)

        # concatenate preprefix (fixed) + prefix (learned) + example
        prefix_ids = prefix_ids[None].to(device).repeat((batch_size, 1)).to(device)
        preprefix_ids = self.preprefix_ids[None].to(device).repeat((batch_size, 1)).to(device)

        if self.prefix_before_input:
            full_input_ids = torch.cat(
                (preprefix_ids, prefix_ids, input_ids), dim=1
            )
            # TODO: How to compute loss at end though??
        else:
            full_input_ids = torch.cat(
                (preprefix_ids, prefix_ids), dim=1
            )
        outputs = torch.cat(
            (
                self.token_embedding.forward(preprefix_ids),
                prefix_embedding[None].repeat((batch_size, 1, 1)),
                self.token_embedding.forward(input_ids)
            ), dim=1
        )
        return full_input_ids, outputs
```

Git & Github

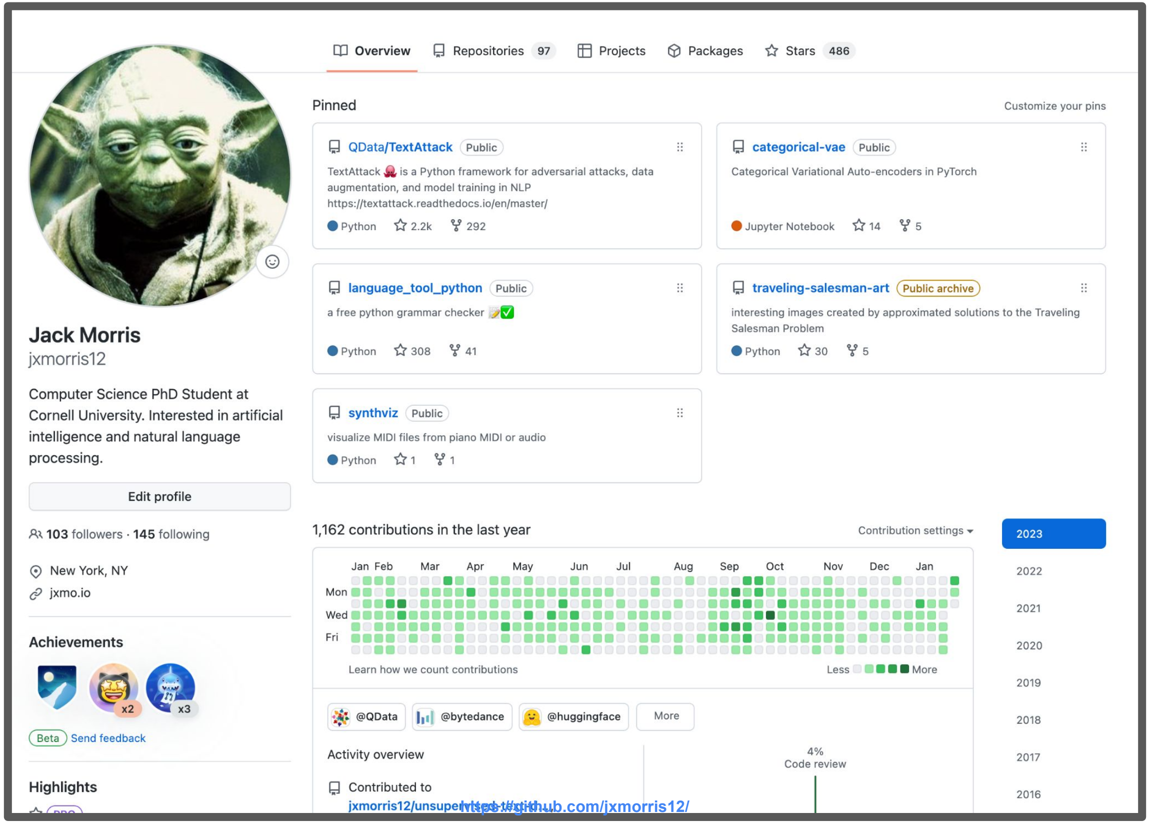

Jack Morris
jxmorris12

Computer Science PhD Student at Cornell University. Interested in artificial intelligence and natural language processing.

103 followers · 145 following

New York, NY
jxmo.io

Achievements

Highlights

Pinned

QData/TextAttack Public
TextAttack 🐙 is a Python framework for adversarial attacks, data augmentation, and model training in NLP
https://textattack.readthedocs.io/en/master/
Python ☆ 2.2k 292

categorical-vae Public
Categorical Variational Auto-encoders in PyTorch
Jupyter Notebook ☆ 14 5

language_tool_python Public
a free python grammar checker 📝✅
Python ☆ 308 41

traveling-salesman-art Public archive
interesting images created by approximated solutions to the Traveling Salesman Problem
Python ☆ 30 5

synthviz Public
visualize MIDI files from piano MIDI or audio
Python ☆ 1 1

1,162 contributions in the last year

Learn how we count contributions

Activity overview

Contributed to
jxmorris12/unsupe

https://github.com/jxmorris12/

Connecting to the cloud

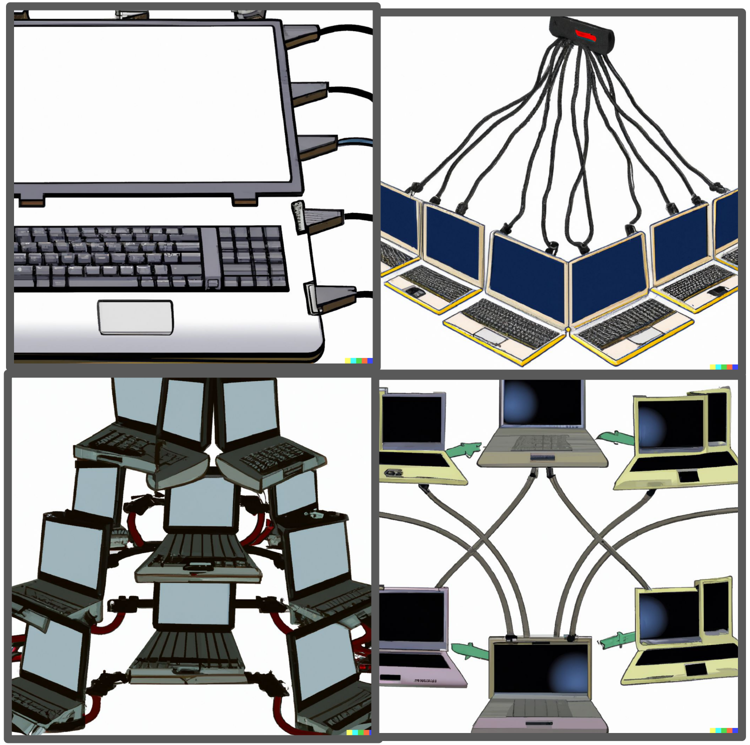

“clip art of a laptop connected by a cord to a large cluster of computers”

Generated by DALL-E 2

Google Colab

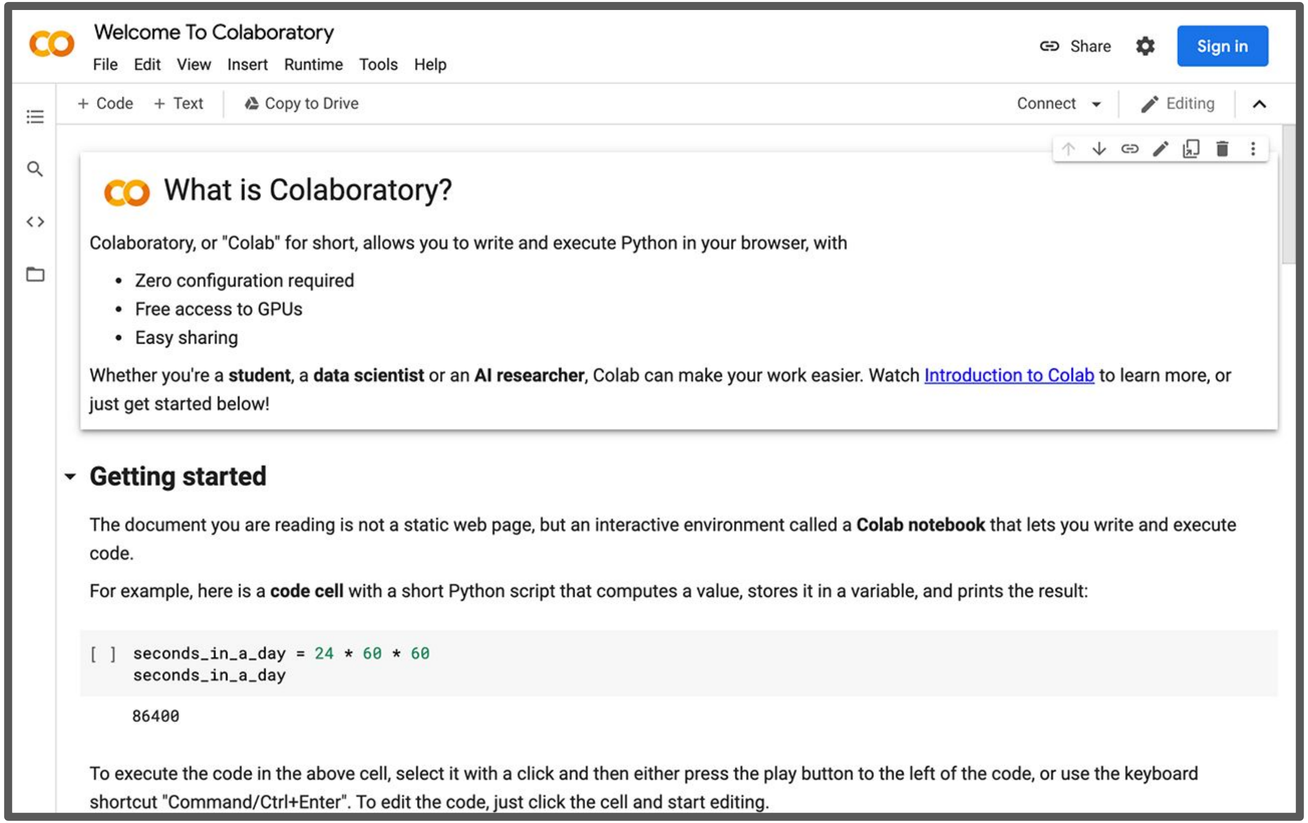

What is Colaboratory?

Colaboratory, or "Colab" for short, allows you to write and execute Python in your browser, with

- Zero configuration required
- Free access to GPUs
- Easy sharing

Whether you're a **student**, a **data scientist** or an **AI researcher**, Colab can make your work easier. Watch Introduction to Colab to learn more, or just get started below!

Getting started

The document you are reading is not a static web page, but an interactive environment called a **Colab notebook** that lets you write and execute code.

For example, here is a **code cell** with a short Python script that computes a value, stores it in a variable, and prints the result:

```
[ ] seconds_in_a_day = 24 * 60 * 60
    seconds_in_a_day
```

```
86400
```

To execute the code in the above cell, select it with a click and then either press the play button to the left of the code, or use the keyboard shortcut "Command/Ctrl+Enter". To edit the code, just click the cell and start editing.

GPUs and CUDA

Anaconda

SciPy
SPYDER
Nun
ASK
bokeh
scikit learn
Numba
Flow
pa
Panel
PYTORCH
matp
CONDA
WS

PyTorch

PyTorch

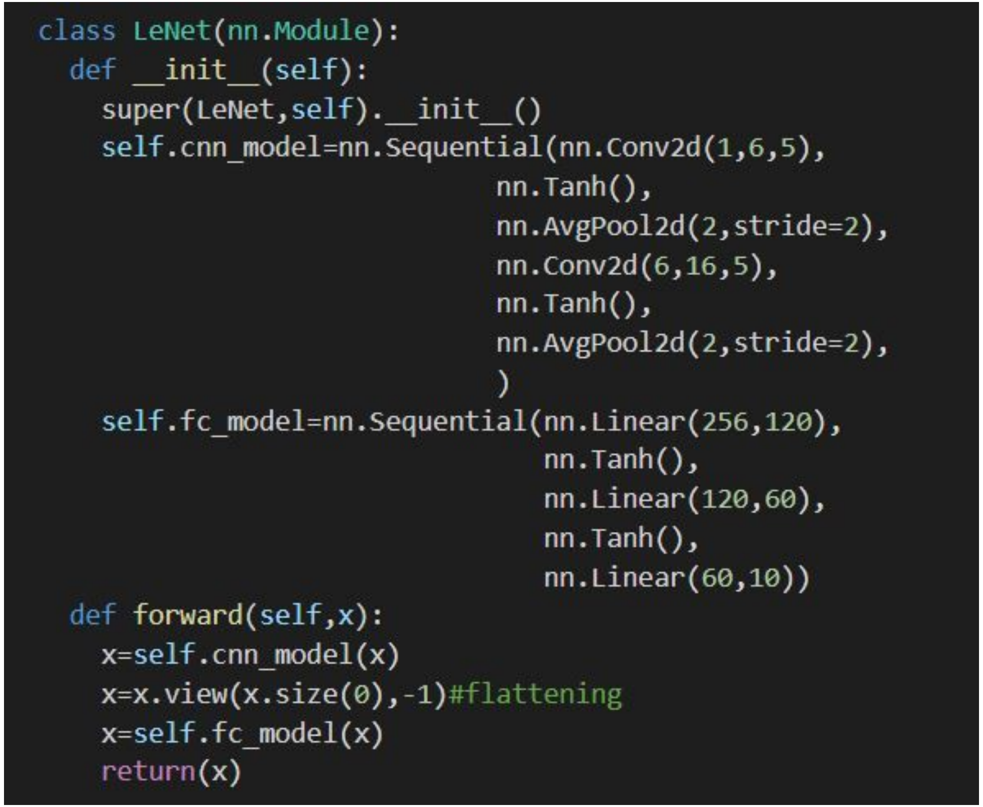

```python
class LeNet(nn.Module):
  def __init__(self):
    super(LeNet,self).__init__()
    self.cnn_model=nn.Sequential(nn.Conv2d(1,6,5),
                              nn.Tanh(),
                              nn.AvgPool2d(2,stride=2),
                              nn.Conv2d(6,16,5),
                              nn.Tanh(),
                              nn.AvgPool2d(2,stride=2),
                              )
    self.fc_model=nn.Sequential(nn.Linear(256,120),
                                 nn.Tanh(),
                                 nn.Linear(120,60),
                                 nn.Tanh(),
                                 nn.Linear(60,10))
  def forward(self,x):
    x=self.cnn_model(x)
    x=x.view(x.size(0),-1)#flattening
    x=self.fc_model(x)
    return(x)
```

HuggingFace (models)

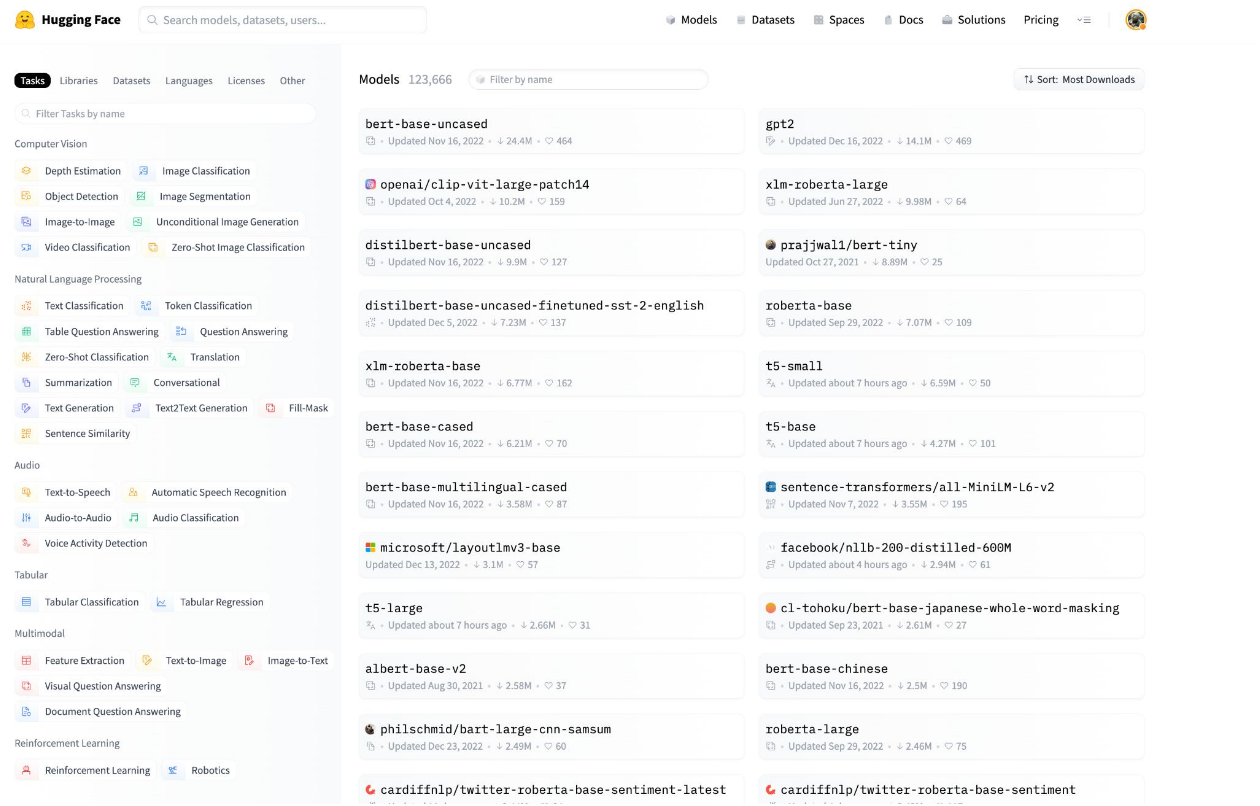

Models 123,666

bert-base-uncased
Updated Nov 16, 2022 • ↓ 24.4M • ♡ 464

gpt2
Updated Dec 16, 2022 • ↓ 14.1M • ♡ 469

openai/clip-vit-large-patch14
Updated Oct 4, 2022 • ↓ 10.2M • ♡ 159

xlm-roberta-large
Updated Jun 27, 2022 • ↓ 9.98M • ♡ 64

distilbert-base-uncased
Updated Nov 16, 2022 • ↓ 9.9M • ♡ 127

prajjwal1/bert-tiny
Updated Oct 27, 2021 • ↓ 8.89M • ♡ 25

distilbert-base-uncased-finetuned-sst-2-english
Updated Dec 5, 2022 • ↓ 7.23M • ♡ 137

roberta-base
Updated Sep 29, 2022 • ↓ 7.07M • ♡ 109

xlm-roberta-base
Updated Nov 16, 2022 • ↓ 6.77M • ♡ 162

t5-small
Updated about 7 hours ago • ↓ 6.59M • ♡ 50

bert-base-cased
Updated Nov 16, 2022 • ↓ 6.21M • ♡ 70

t5-base
Updated about 7 hours ago • ↓ 4.27M • ♡ 101

bert-base-multilingual-cased
Updated Nov 16, 2022 • ↓ 3.58M • ♡ 87

sentence-transformers/all-MiniLM-L6-v2
Updated Nov 7, 2022 • ↓ 3.55M • ♡ 195

microsoft/layoutlmv3-base
Updated Dec 13, 2022 • ↓ 3.1M • ♡ 57

facebook/nllb-200-distilled-600M
Updated about 4 hours ago • ↓ 2.94M • ♡ 61

t5-large
Updated about 7 hours ago • ↓ 2.66M • ♡ 31

cl-tohoku/bert-base-japanese-whole-word-masking
Updated Sep 23, 2021 • ↓ 2.61M • ♡ 27

albert-base-v2
Updated Aug 30, 2021 • ↓ 2.58M • ♡ 37

bert-base-chinese
Updated Nov 16, 2022 • ↓ 2.5M • ♡ 190

philschmid/bart-large-cnn-samsum
Updated Dec 23, 2022 • ↓ 2.49M • ♡ 60

roberta-large
Updated Sep 29, 2022 • ↓ 2.46M • ♡ 75

cardiffnlp/twitter-roberta-base-sentiment-latest

cardiffnlp/twitter-roberta-base-sentiment

HuggingFace (models)

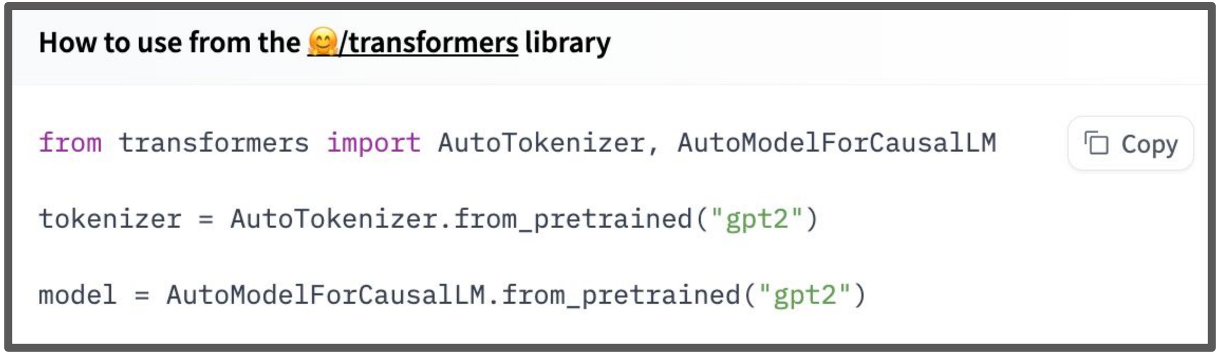

How to use from the 🤗/transformers library

```
from transformers import AutoTokenizer, AutoModelForCausalLM

tokenizer = AutoTokenizer.from_pretrained("gpt2")

model = AutoModelForCausalLM.from_pretrained("gpt2")
```

HuggingFace (datasets)

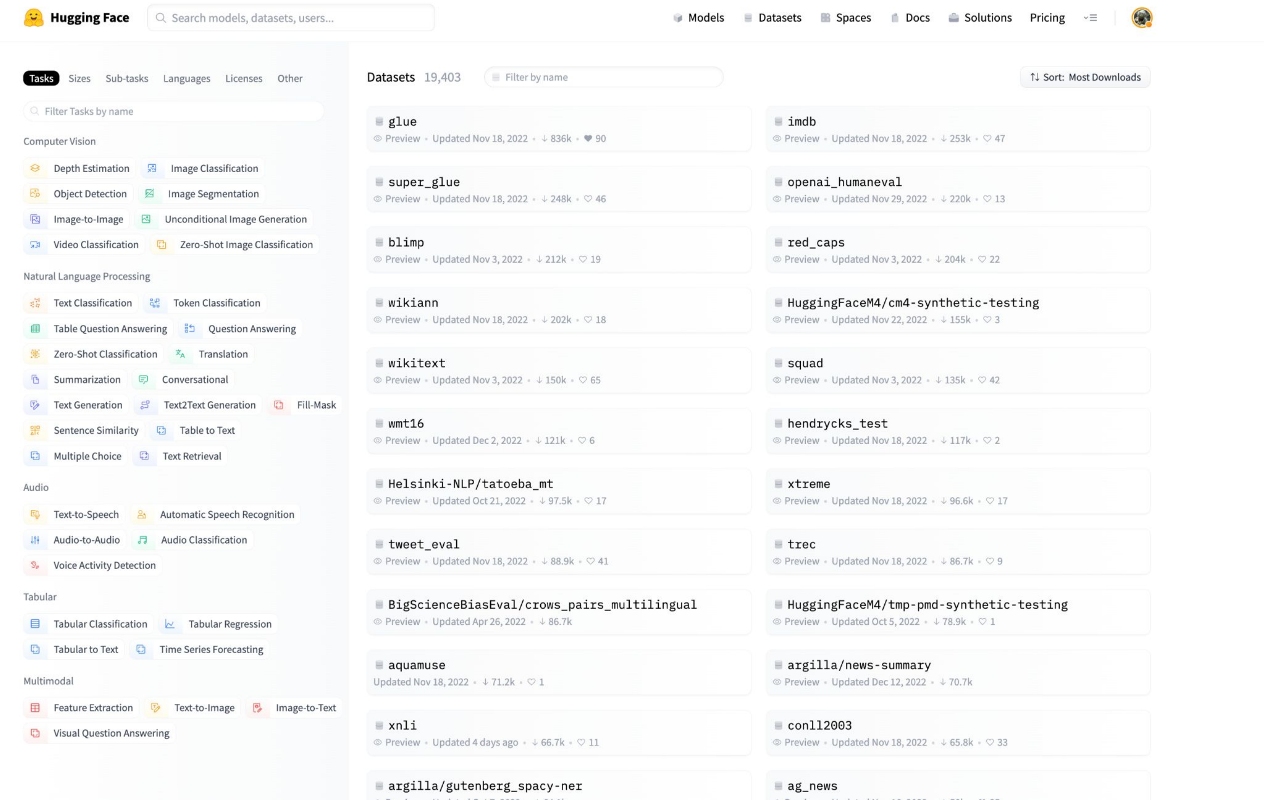

Hugging Face

Search models, datasets, users...

Models Datasets Spaces Docs Solutions Pricing

Tasks Sizes Sub-tasks Languages Licenses Other

Filter Tasks by name

Computer Vision

Depth Estimation, Image Classification, Object Detection, Image Segmentation, Image-to-Image, Unconditional Image Generation, Video Classification, Zero-Shot Image Classification

Natural Language Processing

Text Classification, Token Classification, Table Question Answering, Question Answering, Zero-Shot Classification, Translation, Summarization, Conversational, Text Generation, Text2Text Generation, Fill-Mask, Sentence Similarity, Table to Text, Multiple Choice, Text Retrieval

Audio

Text-to-Speech, Automatic Speech Recognition, Audio-to-Audio, Audio Classification, Voice Activity Detection

Tabular

Tabular Classification, Tabular Regression, Tabular to Text, Time Series Forecasting

Multimodal

Feature Extraction, Text-to-Image, Image-to-Text, Visual Question Answering

Datasets 19,403

Filter by name

Sort: Most Downloads

`glue`
Preview • Updated Nov 18, 2022 • ↓ 836k • ♥ 90

`imdb`
Preview • Updated Nov 18, 2022 • ↓ 253k • ♡ 47

`super_glue`
Preview • Updated Nov 18, 2022 • ↓ 248k • ♡ 46

`openai_humaneval`
Preview • Updated Nov 29, 2022 • ↓ 220k • ♡ 13

`blimp`
Preview • Updated Nov 3, 2022 • ↓ 212k • ♡ 19

`red_caps`
Preview • Updated Nov 3, 2022 • ↓ 204k • ♡ 22

`wikiann`
Preview • Updated Nov 18, 2022 • ↓ 202k • ♡ 18

`HuggingFaceM4/cm4-synthetic-testing`
Preview • Updated Nov 22, 2022 • ↓ 155k • ♡ 3

`wikitext`
Preview • Updated Nov 3, 2022 • ↓ 150k • ♡ 65

`squad`
Preview • Updated Nov 3, 2022 • ↓ 135k • ♡ 42

`wmt16`
Preview • Updated Dec 2, 2022 • ↓ 121k • ♡ 6

`hendrycks_test`
Preview • Updated Nov 18, 2022 • ↓ 117k • ♡ 2

`Helsinki-NLP/tatoeba_mt`
Preview • Updated Oct 21, 2022 • ↓ 97.5k • ♡ 17

`xtreme`
Preview • Updated Nov 18, 2022 • ↓ 96.6k • ♡ 17

`tweet_eval`
Preview • Updated Nov 18, 2022 • ↓ 88.9k • ♡ 41

`trec`
Preview • Updated Nov 18, 2022 • ↓ 86.7k • ♡ 9

`BigScienceBiasEval/crows_pairs_multilingual`
Preview • Updated Apr 26, 2022 • ↓ 86.7k

`HuggingFaceM4/tmp-pmd-synthetic-testing`
Preview • Updated Oct 5, 2022 • ↓ 78.9k • ♡ 1

`aquamuse`
Updated Nov 18, 2022 • ↓ 71.2k • ♡ 1

`argilla/news-summary`
Preview • Updated Dec 12, 2022 • ↓ 70.7k

`xnli`
Preview • Updated 4 days ago • ↓ 66.7k • ♡ 11

`conll2003`
Preview • Updated Nov 18, 2022 • ↓ 65.8k • ♡ 33

`argilla/gutenberg_spacy-ner`

`ag_news`

HuggingFace (models)

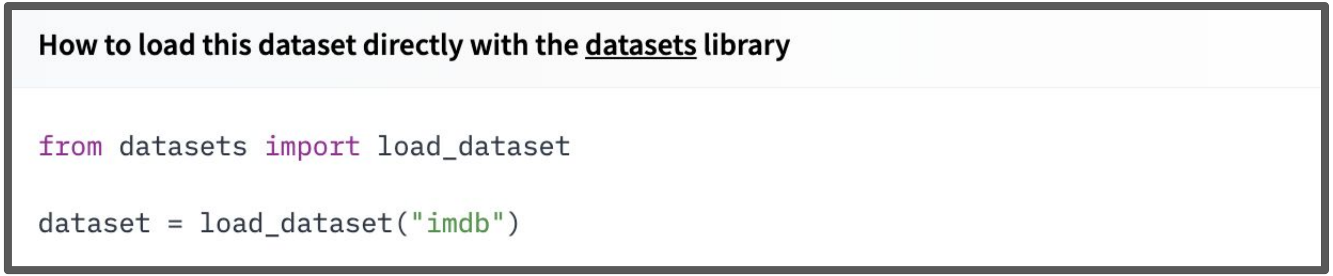

How to load this dataset directly with the datasets library

```
from datasets import load_dataset

dataset = load_dataset("imdb")
```

matplotlib & seaborns

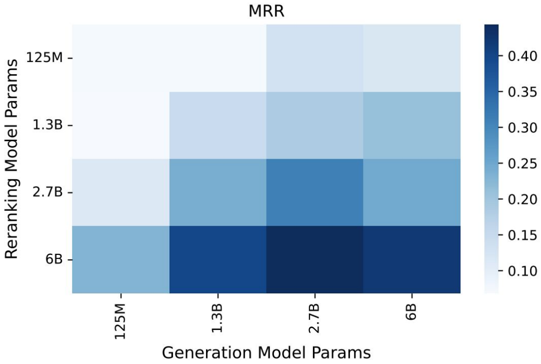

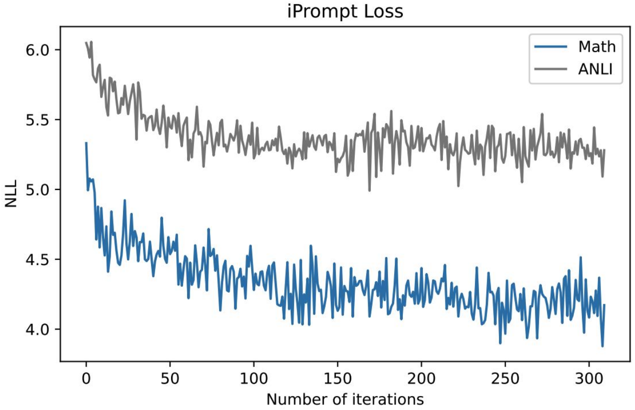

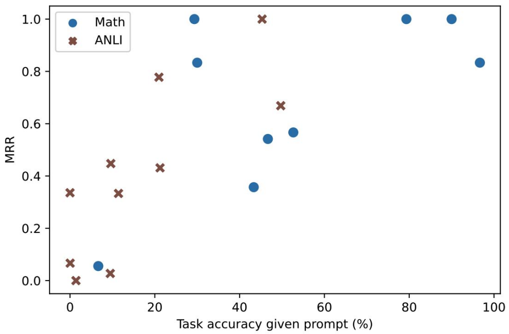

Weights & Biases

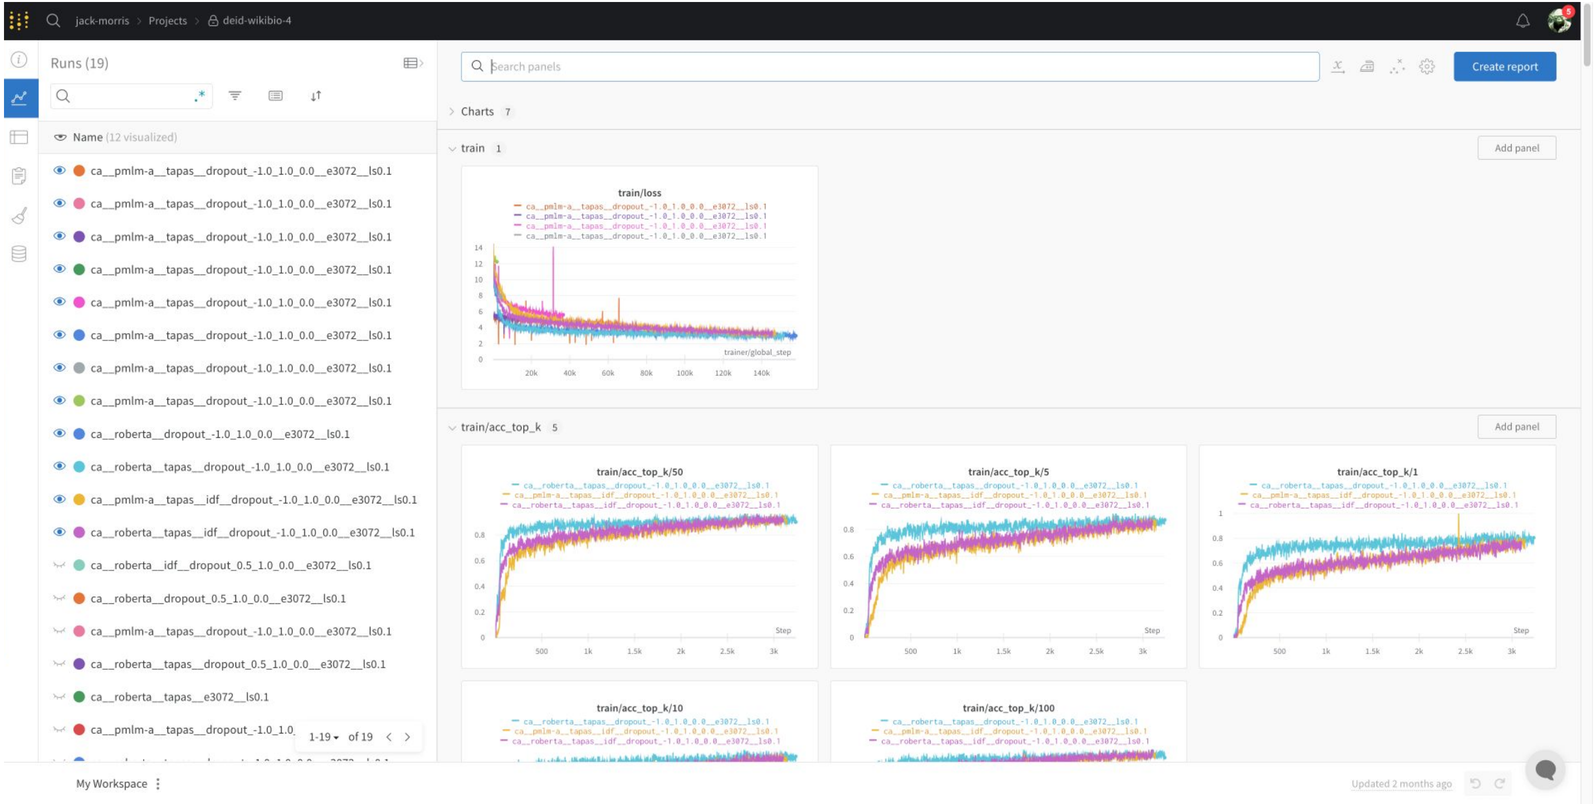

jack-morris > Projects > deid-wikibio-4

Runs (19)

Name (12 visualized)

- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__roberta__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__roberta__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__idf__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__roberta__tapas__idf__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__roberta__idf__dropout_0.5_1.0_0.0__e3072__ls0.1
- ca__roberta__dropout_0.5_1.0_0.0__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0_0.0__e3072__ls0.1
- ca__roberta__tapas__dropout_0.5_1.0_0.0__e3072__ls0.1
- ca__roberta__tapas__e3072__ls0.1
- ca__pmlm-a__tapas__dropout_-1.0_1.0

https://wandb.ai/jack-morris/deid-wikibio-4

Weights & Biases

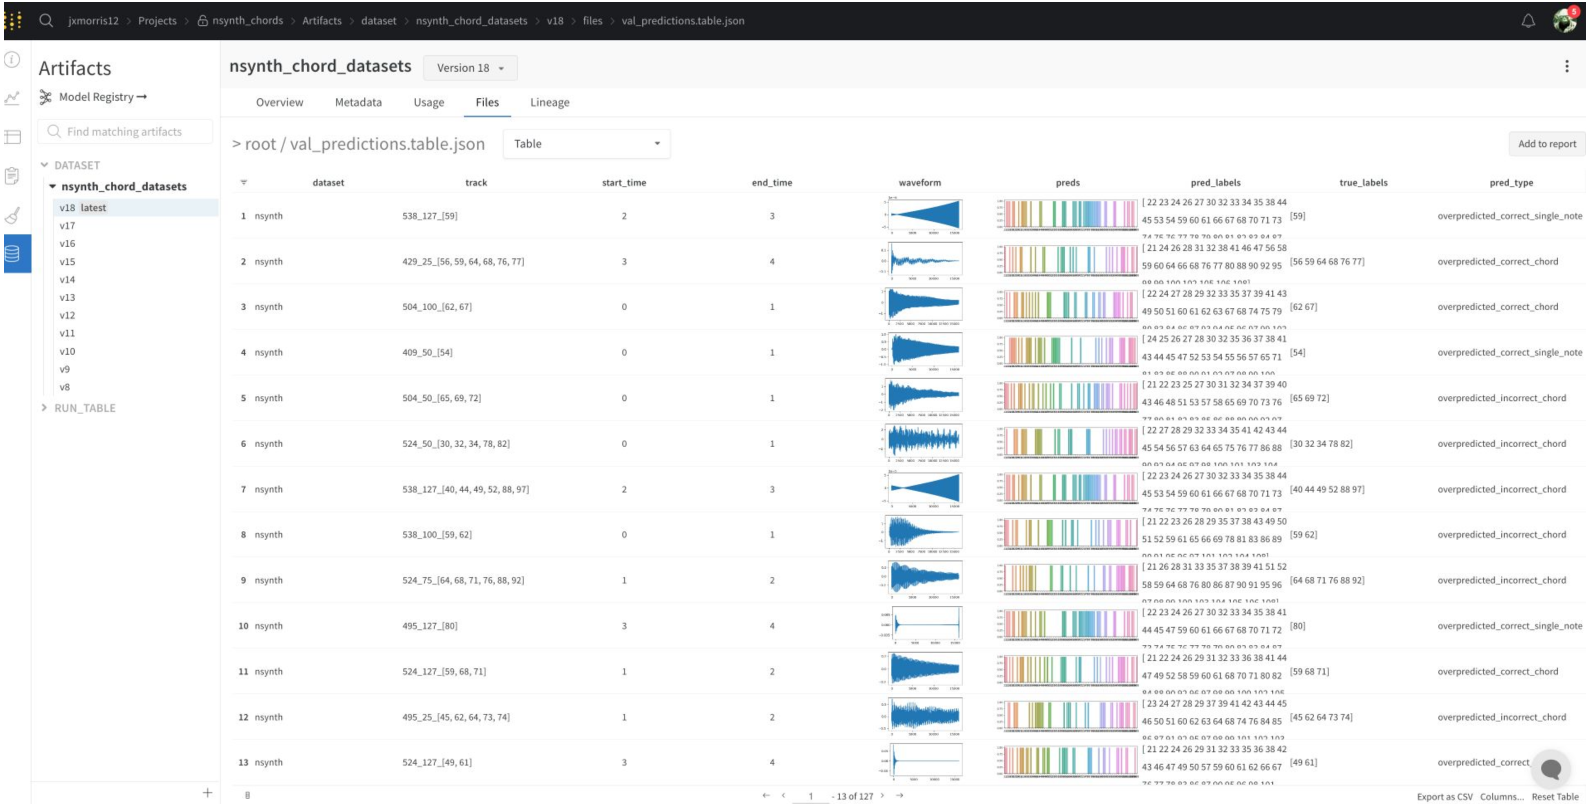

> root / val_predictions.table.json

	dataset	track	start_time	end_time	waveform	preds	pred_labels	true_labels	pred_type
1	nsynth	538_127_[59]	2	3			[22 23 24 26 27 30 32 33 34 35 38 44 45 53 54 59 60 61 66 67 68 70 71 73	[59]	overpredicted_correct_single_note
2	nsynth	429_25_[56, 59, 64, 68, 76, 77]	3	4			[21 24 26 28 31 32 38 41 46 47 56 58 59 60 64 66 68 76 77 80 88 90 92 95	[56 59 64 68 76 77]	overpredicted_correct_chord
3	nsynth	504_100_[62, 67]	0	1			[22 24 27 28 29 32 33 35 37 39 41 43 49 50 51 60 61 62 63 67 68 74 75 79	[62 67]	overpredicted_correct_chord
4	nsynth	409_50_[54]	0	1			[24 25 26 27 28 30 32 35 36 37 38 41 43 44 45 47 52 53 54 55 56 57 65 71	[54]	overpredicted_correct_single_note
5	nsynth	504_50_[65, 69, 72]	0	1			[21 22 23 25 27 30 31 32 34 37 39 40 43 46 48 51 53 57 58 65 69 70 73 76	[65 69 72]	overpredicted_incorrect_chord
6	nsynth	524_50_[30, 32, 34, 78, 82]	0	1			[22 27 28 29 32 33 34 35 41 42 43 44 45 54 56 57 63 64 65 75 76 77 86 88	[30 32 34 78 82]	overpredicted_incorrect_chord
7	nsynth	538_127_[40, 44, 49, 52, 88, 97]	2	3			[22 23 24 26 27 30 32 33 34 35 38 44 45 53 54 59 60 61 66 67 68 70 71 73	[40 44 49 52 88 97]	overpredicted_incorrect_chord
8	nsynth	538_100_[59, 62]	0	1			[21 22 23 26 28 29 35 37 38 43 49 50 51 52 59 61 65 66 69 78 81 83 86 89	[59 62]	overpredicted_incorrect_chord
9	nsynth	524_75_[64, 68, 71, 76, 88, 92]	1	2			[21 26 28 31 33 35 37 38 39 41 51 52 58 59 64 68 76 80 86 87 90 91 95 96	[64 68 71 76 88 92]	overpredicted_incorrect_chord
10	nsynth	495_127_[80]	3	4			[22 23 24 26 27 30 32 33 34 35 38 41 44 45 47 59 60 61 66 67 68 70 71 72	[80]	overpredicted_correct_single_note
11	nsynth	524_127_[59, 68, 71]	1	2			[21 22 24 26 29 31 32 33 36 38 41 44 47 49 52 58 59 60 61 68 70 71 80 82	[59 68 71]	overpredicted_correct_chord
12	nsynth	495_25_[45, 62, 64, 73, 74]	1	2			[23 24 27 28 29 37 39 41 42 43 44 45 46 50 51 60 62 63 64 68 74 76 84 85	[45 62 64 73 74]	overpredicted_incorrect_chord
13	nsynth	524_127_[49, 61]	3	4			[21 22 24 26 29 31 32 33 35 36 38 42 43 46 47 49 50 57 59 60 61 62 66 67	[49 61]	overpredicted_correct

1 - 13 of 127

https://wandb.ai/jxmorris12/nsynth_chords/artifacts/dataset/nsynth_chord_datasets/v18/files/val_predictions.table.json

A deep learning project demo

How do we write a computer program that, given a picture of a car, can predict the make and model? 🚗

Semester overview

Deep learning toolstack ⚒️

Puzzle

Puzzle 🧩

bit.ly/pdl24puzzle1