Weights & Biases

Iterative, Multiplayer Research

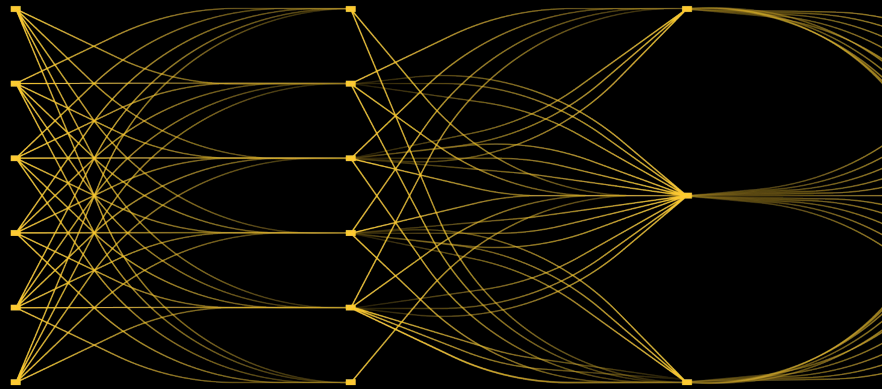

January 2023

Accounts Needed for Pt. 2, Competition

kaggle.com

wandb.ai

colab.research.google.com

The W&B Course

www.wandb.courses

The Goal

Table 3: Detection results on **PASCAL VOC 2007 test set**. The detector is Fast R-CNN and VGG-16. Training data: "07": VOC 2007 trainval, "07+12": union set of VOC 2007 trainval and VOC 2012 trainval. For RPN, the train-time proposals for Fast R-CNN are 2000. †: this number was reported in [2]; using the repository provided by this paper, this result is higher (68.1).

method	# proposals	data	mAP (%)
SS	2000	07	66.9†
SS	2000	07+12	70.0
RPN+VGG, unshared	300	07	68.5
RPN+VGG, shared	300	07	69.9
RPN+VGG, shared	300	07+12	**73.2**
RPN+VGG, shared	300	COCO+07+12	**78.8**

Table 4: Detection results on **PASCAL VOC 2012 test set**. The detector is Fast R-CNN and VGG-16. Training data: "07": VOC 2007 trainval, "07++12": union set of VOC 2007 trainval+test and VOC 2012 trainval. For RPN, the train-time proposals for Fast R-CNN are 2000. †: http://host.robots.ox.ac.uk:8080/anonymous/HZJTQA.html. ‡: http://host.robots.ox.ac.uk:8080/anonymous/YNPLXB.html. §: http://host.robots.ox.ac.uk:8080/anonymous/XEDH10.html.

method	# proposals	data	mAP (%)
SS	2000	12	65.7
SS	2000	07++12	68.4
RPN+VGG, shared†	300	12	67.0
RPN+VGG, shared‡	300	07++12	**70.4**
RPN+VGG, shared§	300	COCO+07++12	**75.9**

Table 5: **Timing** (ms) on a K40 GPU, except SS proposal is evaluated in a CPU. "Region-wise" includes NMS, pooling, fully-connected, and softmax layers. See our released code for the profiling of running time.

model	system	conv	proposal	region-wise	total	rate
VGG	SS + Fast R-CNN	146	1510	174	1830	0.5 fps
VGG	RPN + Fast R-CNN	141	**10**	47	**198**	**5 fps**
ZF	RPN + Fast R-CNN	31	**3**	25	**59**	**17 fps**

	backbone	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Two-stage methods							
Faster R-CNN+++ [5]	ResNet-101-C4	34.9	55.7	37.4	15.6	38.7	50.9
Faster R-CNN w FPN [8]	ResNet-101-FPN	36.2	59.1	39.0	18.2	39.0	48.2
Faster R-CNN by G-RMI [6]	Inception-ResNet-v2 [21]	34.7	55.5	36.7	13.5	38.1	52.0
Faster R-CNN w TDM [20]	Inception-ResNet-v2-TDM	36.8	57.7	39.2	16.2	39.8	**52.1**
One-stage methods							
YOLOv2 [15]	DarkNet-19 [15]	21.6	44.0	19.2	5.0	22.4	35.5
SSD513 [11, 3]	ResNet-101-SSD	31.2	50.4	33.3	10.2	34.5	49.8
DSSD513 [3]	ResNet-101-DSSD	33.2	53.3	35.2	13.0	35.4	51.1
RetinaNet [9]	ResNet-101-FPN	39.1	59.1	42.3	21.8	42.7	50.2
RetinaNet [9]	ResNeXt-101-FPN	**40.8**	**61.1**	**44.1**	**24.1**	**44.2**	51.2
YOLOv3 608 × 608	Darknet-53	33.0	57.9	34.4	18.3	35.4	41.9

0000 p1 DRZ 3.11 - not super eval
mag_threshold 0.21
2L 300N BLSTM (BasicLSTM)
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(x+1.0)
103300 training, 2000 CV
model:
weights20170224-005946_v10.1419
(p1, loss .1419, epoch 40 [task0])
MEAN IBM SDR GAIN: 2.324 -
with 0.15 thresh during cluster
STD IBM SDR GAIN: 2.276
MEAN IBM SDR GAIN: 2.110 -
with 0.32 threshold during cluster
STD IBM SDR GAIN: 2.254

0001 p2 DRZ 3.23
mag_threshold 0.12
2L 300N BLSTM_clean (LSTM & many reworks) - note, this was the massive model rewrite
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(x+1.0)
10330 training, 2000 CV
model:
weights20170224-032054_v10.1418
(p2, loss .1418, epoch 40 [task0])
MEAN IBM SDR GAIN: 2.056 -
with 0.15 thresh during cluster
STD IBM SDR GAIN: 2.214
MEAN IBM SDR GAIN: 2.068 -
with 0.32 threshold during cluster
STD IBM SDR GAIN: 2.205
MEAN IBM SDR GAIN: 2.108 - with fancy best SDR of the two system (eval_sdr2.py)
STD IBM SDR GAIN: 4.114
MEAN IBM SDR GAIN: 5.915 - using EXACT script measuring both voices gain (SUPER_EVAL)
STD IBM SDR GAIN: 4.349

0000 p1 DRZ 3.11 - not super eval
mag_threshold 0.21
2L 300N BLSTM (BasicLSTM)
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(x+1.0)
103300 training, 2000 CV
model:
weights20170224-005946_v10
(p1, loss .1419, epoch 40
MEAN IBM SDR GAIN: 2.
with 0.15 thresh during cl
STD IBM SDR GAIN: 2.
MEAN IBM SDR GAIN: 2.
with 0.32 threshold during
STD IBM SDR GAIN: 2.

0001 p2 DRZ 3.23
mag_threshold 0.12
2L 300N BLSTM_clean (LSTM & many reworks) - note, this was the massive model rewrite
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label

A	B	L	M	N	O	P
Experiment Name	Created	train_loss	valid_loss	acc	traffic_acc	road_acc
best car acc (50% data)	2021-04-14	0.5375041962	0.442730248	0.8823291659	0.8663836718	0.9399003386
best traffic acc (50% data))	2021-04-14	0.4919361174	0.4202951491	0.8879730701	0.8718349934	0.9439761043
best overall IOU (20% data)	2021-04-14	0.5095784068	0.4658596516	0.8725891709	0.8592621684	0.9359762073
major-sweep-196	2021-01-31	0.5705417991	0.4875227213	0.8698127866	0.8570468426	0.9454026222
swept-sweep-164	2021-01-31	0.5535062551	0.4849829972	0.8701210618	0.8567070365	0.9204238057
silver-sweep-139	2021-01-31	0.563354373	0.5251165628	0.871628046	0.846842885	0.9262287617
laced-sweep-115	2021-01-31	0.5277443528	0.5124291778	0.8705932498	0.8521561027	0.9389513731
eager-sweep-97	2021-01-31	0.5488699675	0.5005864501	0.8738754392	0.8612990975	0.913561523
rich-sweep-88	2021-01-31	0.5587444901	0.5211353302	0.8785927892	0.8512274623	0.9295567274
hopeful-sweep-33	2021-01-31	0.503461957	0.4650281966	0.8706912994	0.8560319543	0.9387732744
autumn-sweep-24	2021-01-31	0.5777919888	0.500880897	0.8755427003	0.8561192751	0.9181208611
decent-sweep-21	2021-01-31	0.5714729428	0.4979581237	0.8745227456	0.8490597606	0.9463140965
vague-sweep-5	2021-01-31	0.6230331063	0.473508656	0.8732874393	0.8601382971	0.9297611117
Second best acc	2021-01-31	0.4194990396	0.4509823024	0.8873019218	0.8705806732	0.9445936084

0000 p1 DRZ 3.11 - not super eval
mag_threshold 0.21
2L 300N BLSTM (BasicLSTM)
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(x+1.0)
103300 training, 2000 CV
model:
weights20170224-005946_v10
(p1, loss .1419, epoch 40
MEAN IBM SDR GAIN: 2.
with 0.15 thresh during cl
STD IBM SDR GAIN: 2.
MEAN IBM SDR GAIN: 2.
with 0.32 threshold during
STD IBM SDR GAIN: 2.

0001 p2 DRZ 3.23
mag_threshold 0.12
2L 300N BLSTM_clean (LSTM & many reworks) - note, this was the massive model rewrite
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label

A	B	L	M	N	O	P
Experiment Name	Created	train_loss	valid_loss	acc	traffic_acc	road_acc
best car acc (50% data)	2021-04-14	0.5375041962	0.442730248	0.8823291659	0.8663836718	0.9399003386
best traffic acc (50% data))	2021-04-14	0.4919361174	0.4202951491	0.8879730701	0.8718349934	0.9439761043
best overall IOU (20% data)	2021-04-14	0.5095784068	0.4658596516	0.8725891709	0.8592621684	0.9359762073
major-sweep-196	2021-01-31	0.5705417991	0.4875227213			
swept-sweep-164	2021-01-31	0.5535062551	0.4849829972			
silver-sweep-139	2021-01-31	0.563354373	0.5251165628			
laced-sweep-115	2021-01-31	0.5277443528	0.5124291778			
eager-sweep-97	2021-01-31	0.5488699675	0.5005864501			
rich-sweep-88	2021-01-31	0.5587444901	0.5211353302			
hopeful-sweep-33	2021-01-31	0.503461957	0.4650281966			
autumn-sweep-24	2021-01-31	0.5777919888	0.500880897			
decent-sweep-21	2021-01-31	0.5714729428	0.4979581237			
vague-sweep-5	2021-01-31	0.6230331063	0.473508656			
Second best acc	2021-01-31	0.4194990396	0.4509823024			

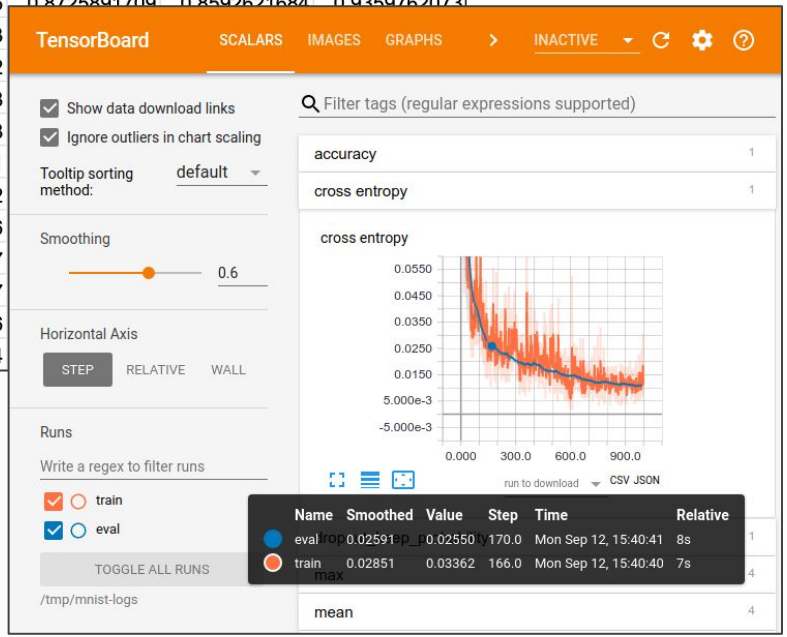

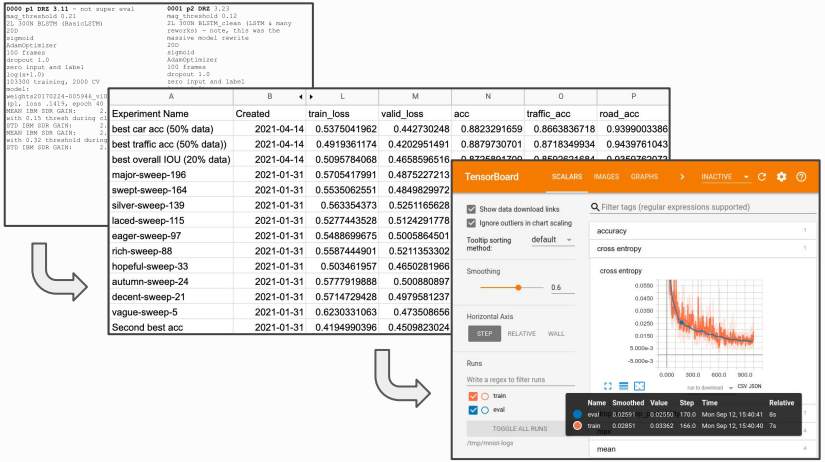
0000 p1 DRZ 3.11 - not super eval
mag_threshold 0.21
2L 300N BLSTM (basicLSTM)
200
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(s+1.0)
103300 training, 2000 CV
model:
0001 p2 DRZ 3.23
mag_threshold 0.12
2L 300N BLSTM_clean (LSTM & many reworks) - note, this was the massive model rewrite
200
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
Experiment Name
Created
train_loss
valid_loss
acc
traffic_acc
road_acc
best car acc (50% data) 2021-04-14 0.5375041962 0.442730248 0.8823291659 0.8663836718 0.9399003386
best traffic acc (50% data)) 2021-04-14 0.4919361174 0.4202951491 0.8879730701 0.8718349934 0.9439761043
best overall IOU (20% data) 2021-04-14 0.5095784068 0.4658596516
major-sweep-196 2021-01-31 0.5705417991 0.4875227213
swept-sweep-164 2021-01-31 0.5535062551 0.4849829972
silver-sweep-139 2021-01-31 0.563354373 0.5251165628
laced-sweep-115 2021-01-31 0.5277443528 0.5124291778
eager-sweep-97 2021-01-31 0.5488699675 0.5005864501
rich-sweep-88 2021-01-31 0.5587444901 0.5211353302
hopeful-sweep-33 2021-01-31 0.503461957 0.4650281966
autumn-sweep-24 2021-01-31 0.5777919888 0.500880897
decent-sweep-21 2021-01-31 0.5714729428 0.4979581237
vague-sweep-5 2021-01-31 0.6230331063 0.473508656
Second best acc 2021-01-31 0.4194990396 0.4509823024
TensorBoard
SCALARS
IMAGES
GRAPHS
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method:
default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
train
eval
TOGGLE ALL RUNS
/tmp/mnist-logs
Filter tags (regular expressions supported)
accuracy
cross entropy
mean
Name Smoothed Value Step Time Relative
eval 0.02591 0.02550 170.0 Mon Sep 12, 15:40:41 8s
train 0.02851 0.03362 166.0 Mon Sep 12, 15:40:40 7s

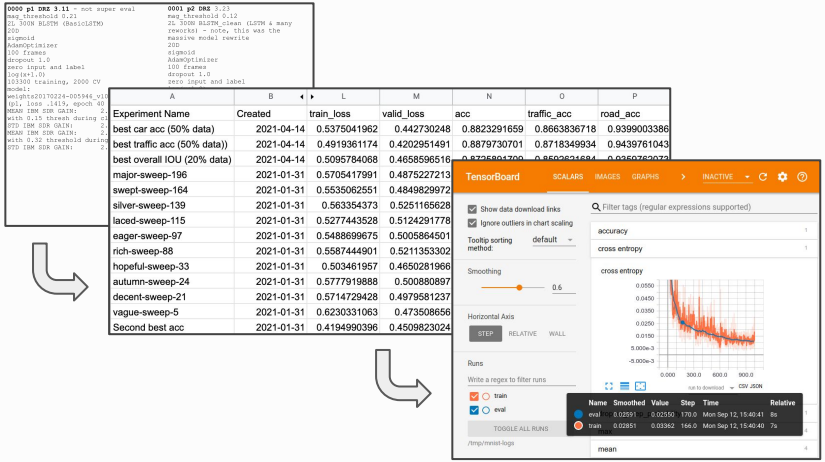
0000 p1 DRZ 3.11 - not super eval
mag_threshold 0.21
2L 300N BLSTM (basicLSTM)
200
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(s+1.0)
103300 training, 2000 CV
model:
0001 p2 DRZ 3.23
mag_threshold 0.12
2L 300N BLSTM_clean (LSTM & many reworks) - note, this was the massive model rewrite
200
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
Experiment Name
Created
train_loss
valid_loss
acc
traffic_acc
road_acc
best car acc (50% data) 2021-04-14 0.5375041962 0.442730248 0.8823291659 0.8663836718 0.9399003386
best traffic acc (50% data)) 2021-04-14 0.4919361174 0.4202951491 0.8879730701 0.8718349934 0.9439761043
best overall IOU (20% data) 2021-04-14 0.5095784068 0.4658596516
major-sweep-196 2021-01-31 0.5705417991 0.4875227213
swept-sweep-164 2021-01-31 0.5535062551 0.4849829972
silver-sweep-139 2021-01-31 0.563354373 0.5251165628
laced-sweep-115 2021-01-31 0.5277443528 0.5124291778
eager-sweep-97 2021-01-31 0.5488699675 0.5005864501
rich-sweep-88 2021-01-31 0.5587444901 0.5211353302
hopeful-sweep-33 2021-01-31 0.503461957 0.4650281966
autumn-sweep-24 2021-01-31 0.5777919888 0.500880897
decent-sweep-21 2021-01-31 0.5714729428 0.4979581237
vague-sweep-5 2021-01-31 0.6230331063 0.473508656
Second best acc 2021-01-31 0.4194990396 0.4509823024
TensorBoard
SCALARS
IMAGES
GRAPHS
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method:
default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
train
eval
TOGGLE ALL RUNS
/tmp/mnist-logs
Filter tags (regular expressions supported)
accuracy
cross entropy
mean
Name Smoothed Value Step Time Relative
eval 0.02591 0.02550 170.0 Mon Sep 12, 15:40:41 8s
train 0.02851 0.03362 166.0 Mon Sep 12, 15:40:40 7s

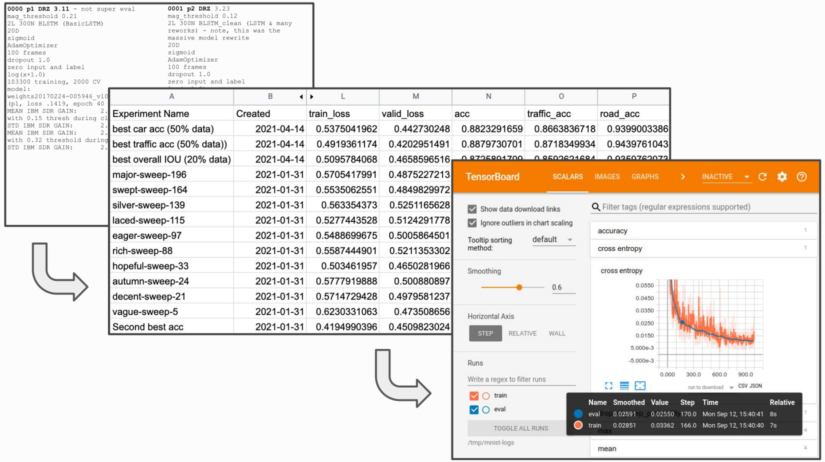
0000 p1 DRZ 3.11 - not super eval
mag_threshold 0.21
2L 300N BLSTM (BasicLSTM)
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
log(x+1.0)
103300 training, 2000 CV
model:
0001 p2 DRZ 3.23
mag_threshold 0.12
2L 300N BLSTM_clean (LSTM & many reworks) - note, this was the massive model rewrite
20D
sigmoid
AdamOptimizer
100 frames
dropout 1.0
zero input and label
Experiment Name | Created | train_loss | valid_loss | acc | traffic_acc | road_acc
best car acc (50% data) | 2021-04-14 | 0.5375041962 | 0.442730248 | 0.8823291659 | 0.8663836718 | 0.9399003386
best traffic acc (50% data)) | 2021-04-14 | 0.4919361174 | 0.4202951491 | 0.8879730701 | 0.8718349934 | 0.9439761043
best overall IOU (20% data) | 2021-04-14 | 0.5095784068 | 0.4658596516
major-sweep-196 | 2021-01-31 | 0.5705417991 | 0.4875227213
swept-sweep-164 | 2021-01-31 | 0.5535062551 | 0.4849829972
silver-sweep-139 | 2021-01-31 | 0.563354373 | 0.5251165628
laced-sweep-115 | 2021-01-31 | 0.5277443528 | 0.5124291778
eager-sweep-97 | 2021-01-31 | 0.5488699675 | 0.5005864501
rich-sweep-88 | 2021-01-31 | 0.5587444901 | 0.5211353302
hopeful-sweep-33 | 2021-01-31 | 0.503461957 | 0.4650281966
autumn-sweep-24 | 2021-01-31 | 0.5777919888 | 0.500880897
decent-sweep-21 | 2021-01-31 | 0.5714729428 | 0.4979581237
vague-sweep-5 | 2021-01-31 | 0.6230331063 | 0.473508656
Second best acc | 2021-01-31 | 0.4194990396 | 0.4509823024
TensorBoard
SCALARS
IMAGES
GRAPHS
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method: default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
train
eval
TOGGLE ALL RUNS
/tmp/mnist-logs
Filter tags (regular expressions supported)
accuracy
cross entropy
cross entropy
Name | Smoothed | Value | Step | Time | Relative
eval | 0.02591 | 0.02550 | 170.0 | Mon Sep 12, 15:40:41 | 8s
train | 0.02851 | 0.03362 | 166.0 | Mon Sep 12, 15:40:40 | 7s
mean

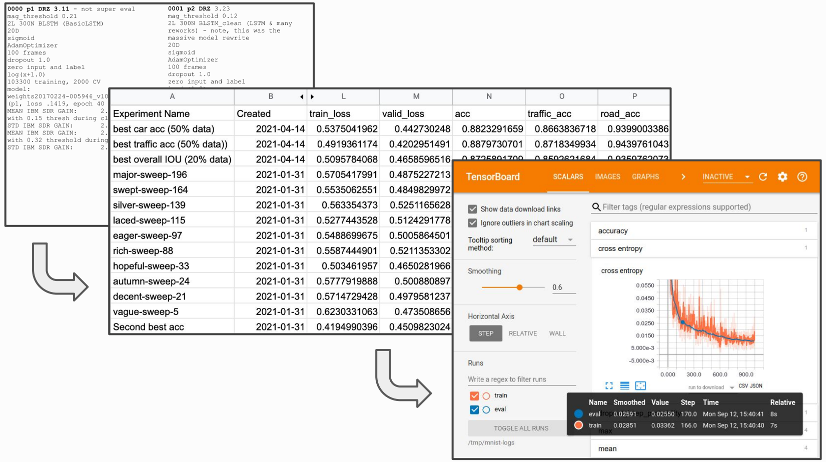

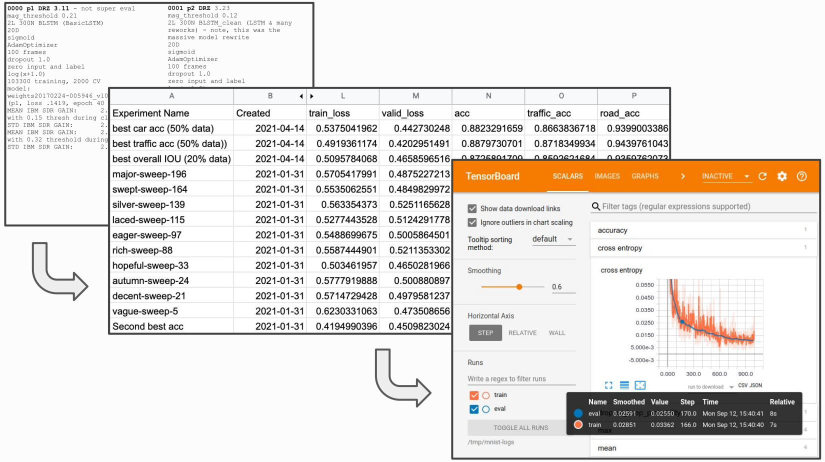

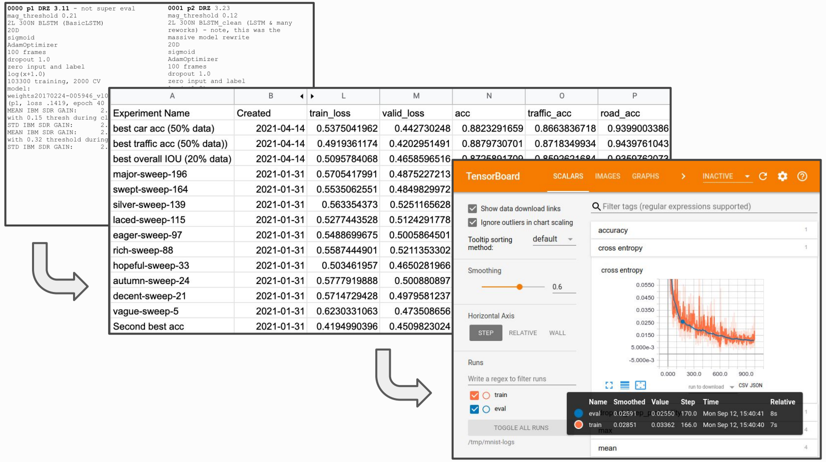

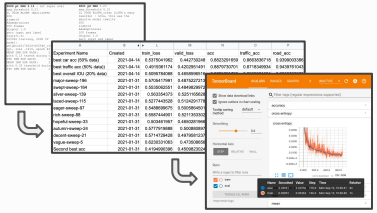
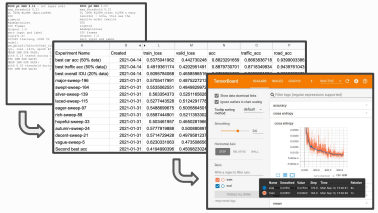
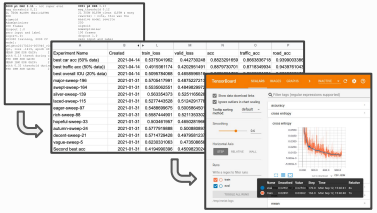
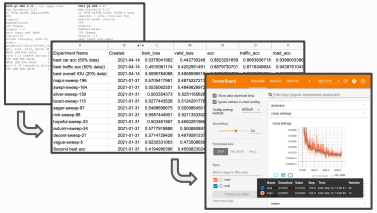
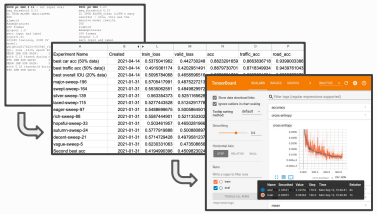
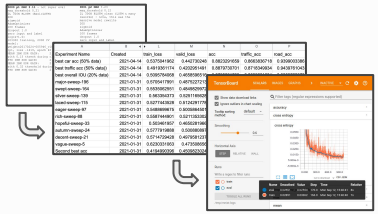
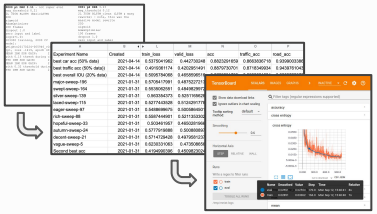
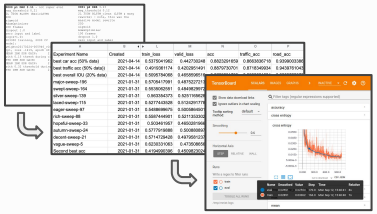
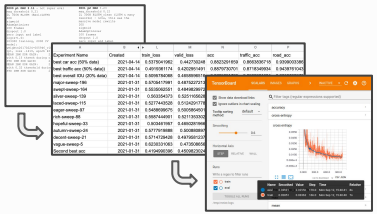
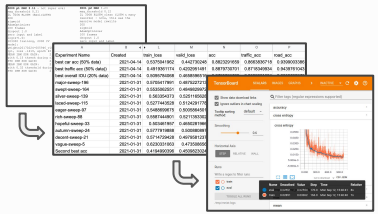
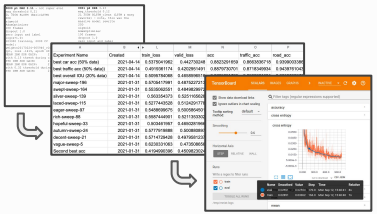
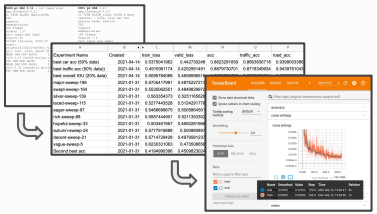
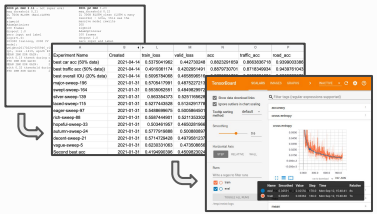
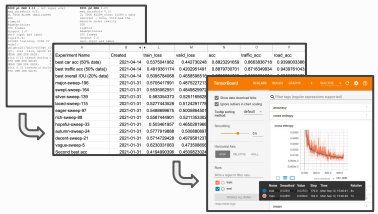
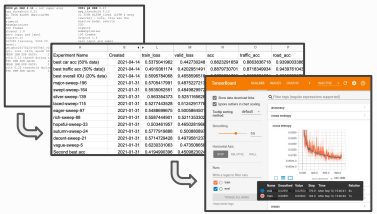
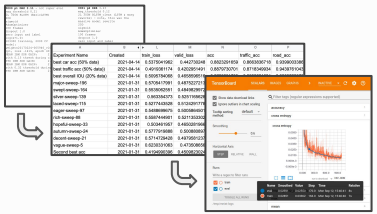

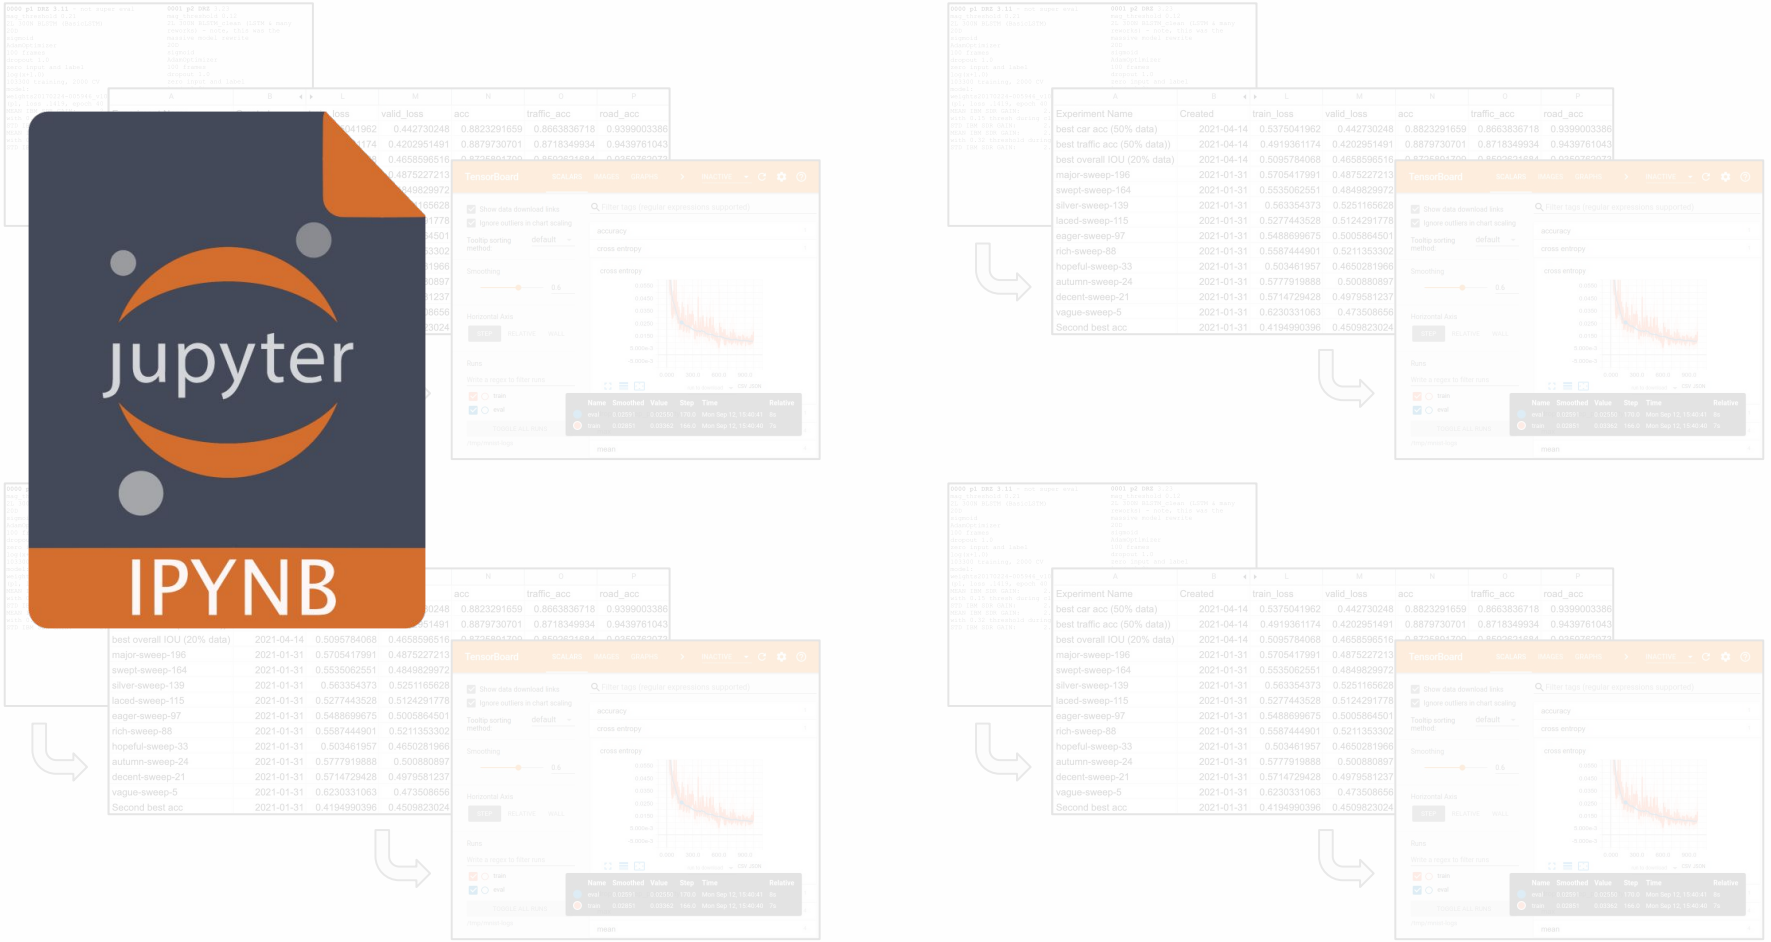
jupyter
IPYNB

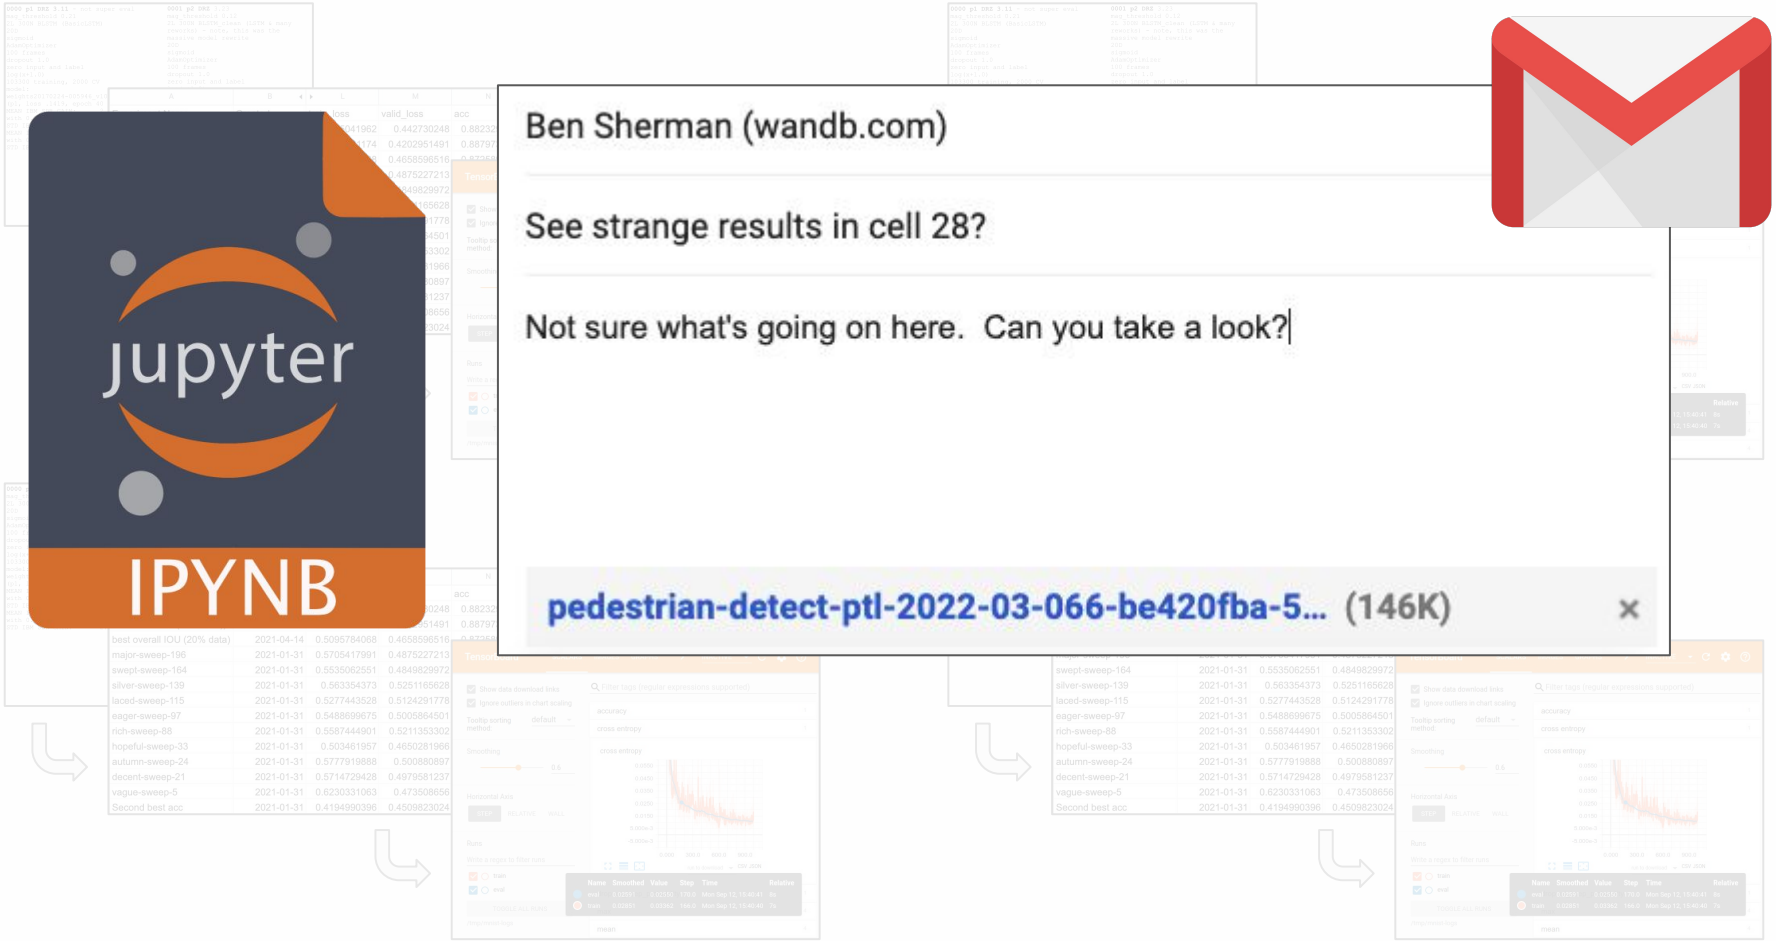
jupyter
IPYNB
Ben Sherman (wandb.com)
See strange results in cell 28?
Not sure what's going on here. Can you take a look?
pedestrian-detect-ptl-2022-03-066-be420fba-5... (146K)

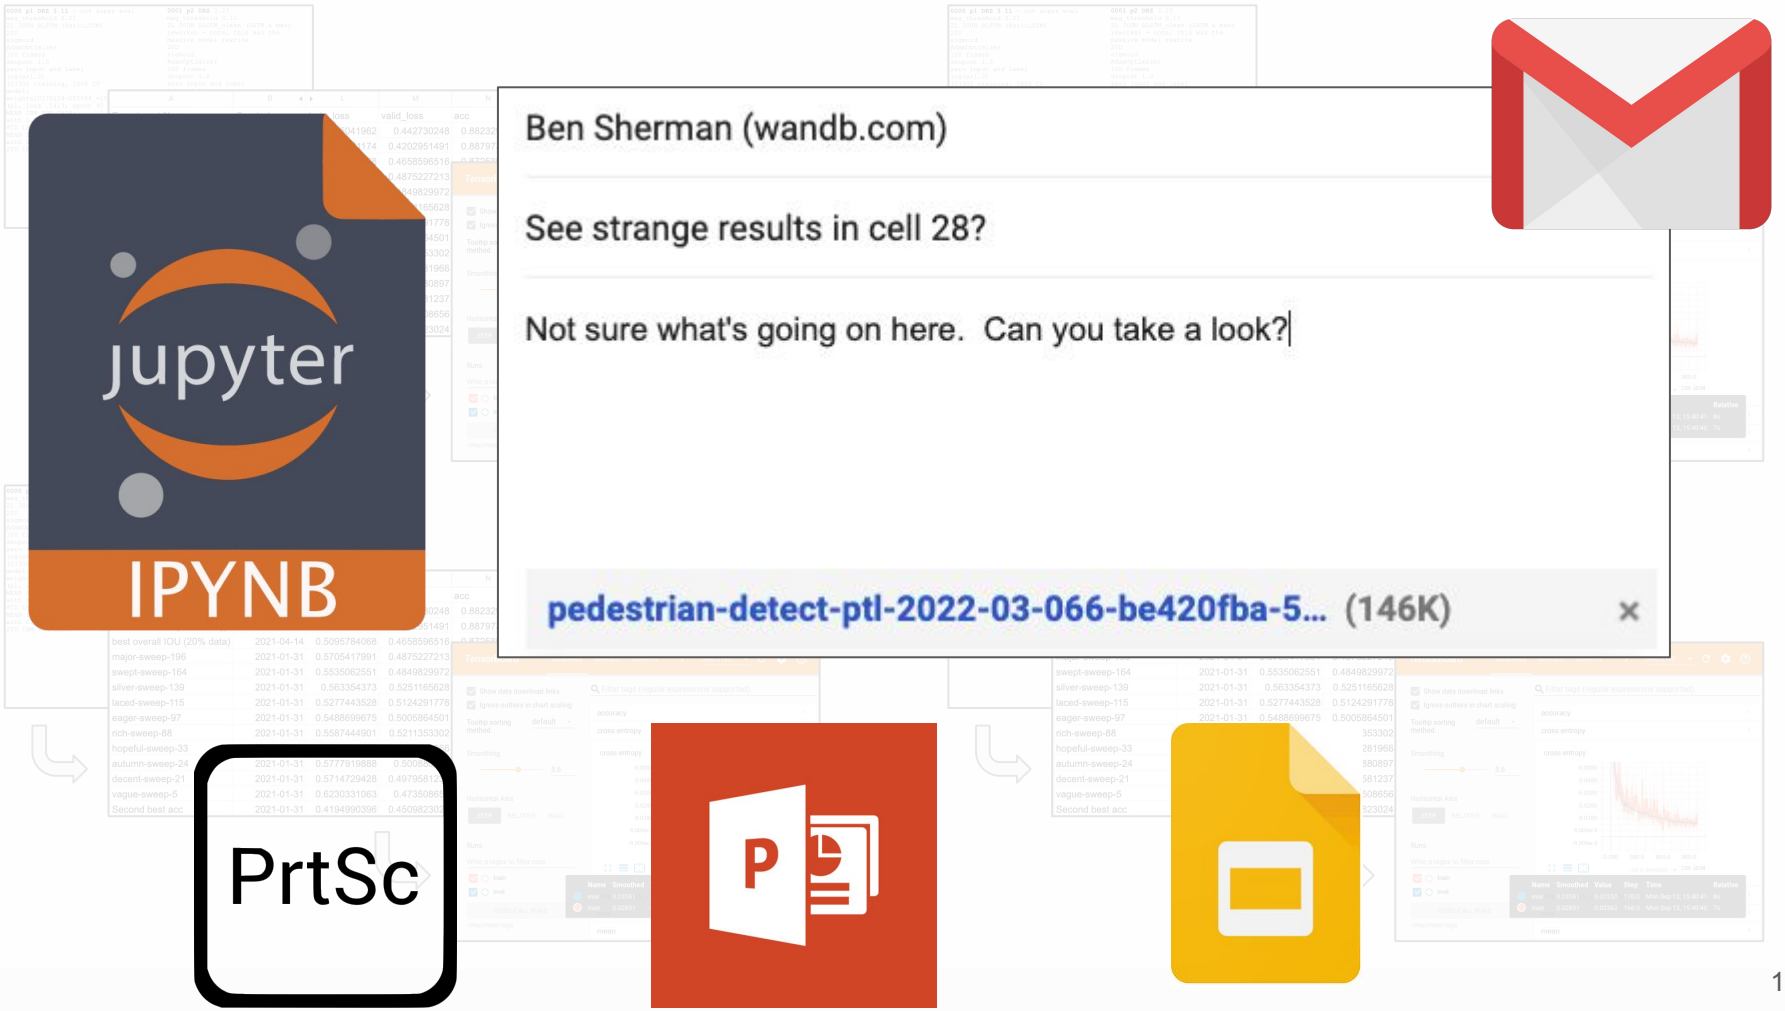
jupyter
IPYNB
Ben Sherman (wandb.com)
See strange results in cell 28?
Not sure what's going on here. Can you take a look?
pedestrian-detect-ptl-2022-03-066-be420fba-5... (146K)
PrtSc

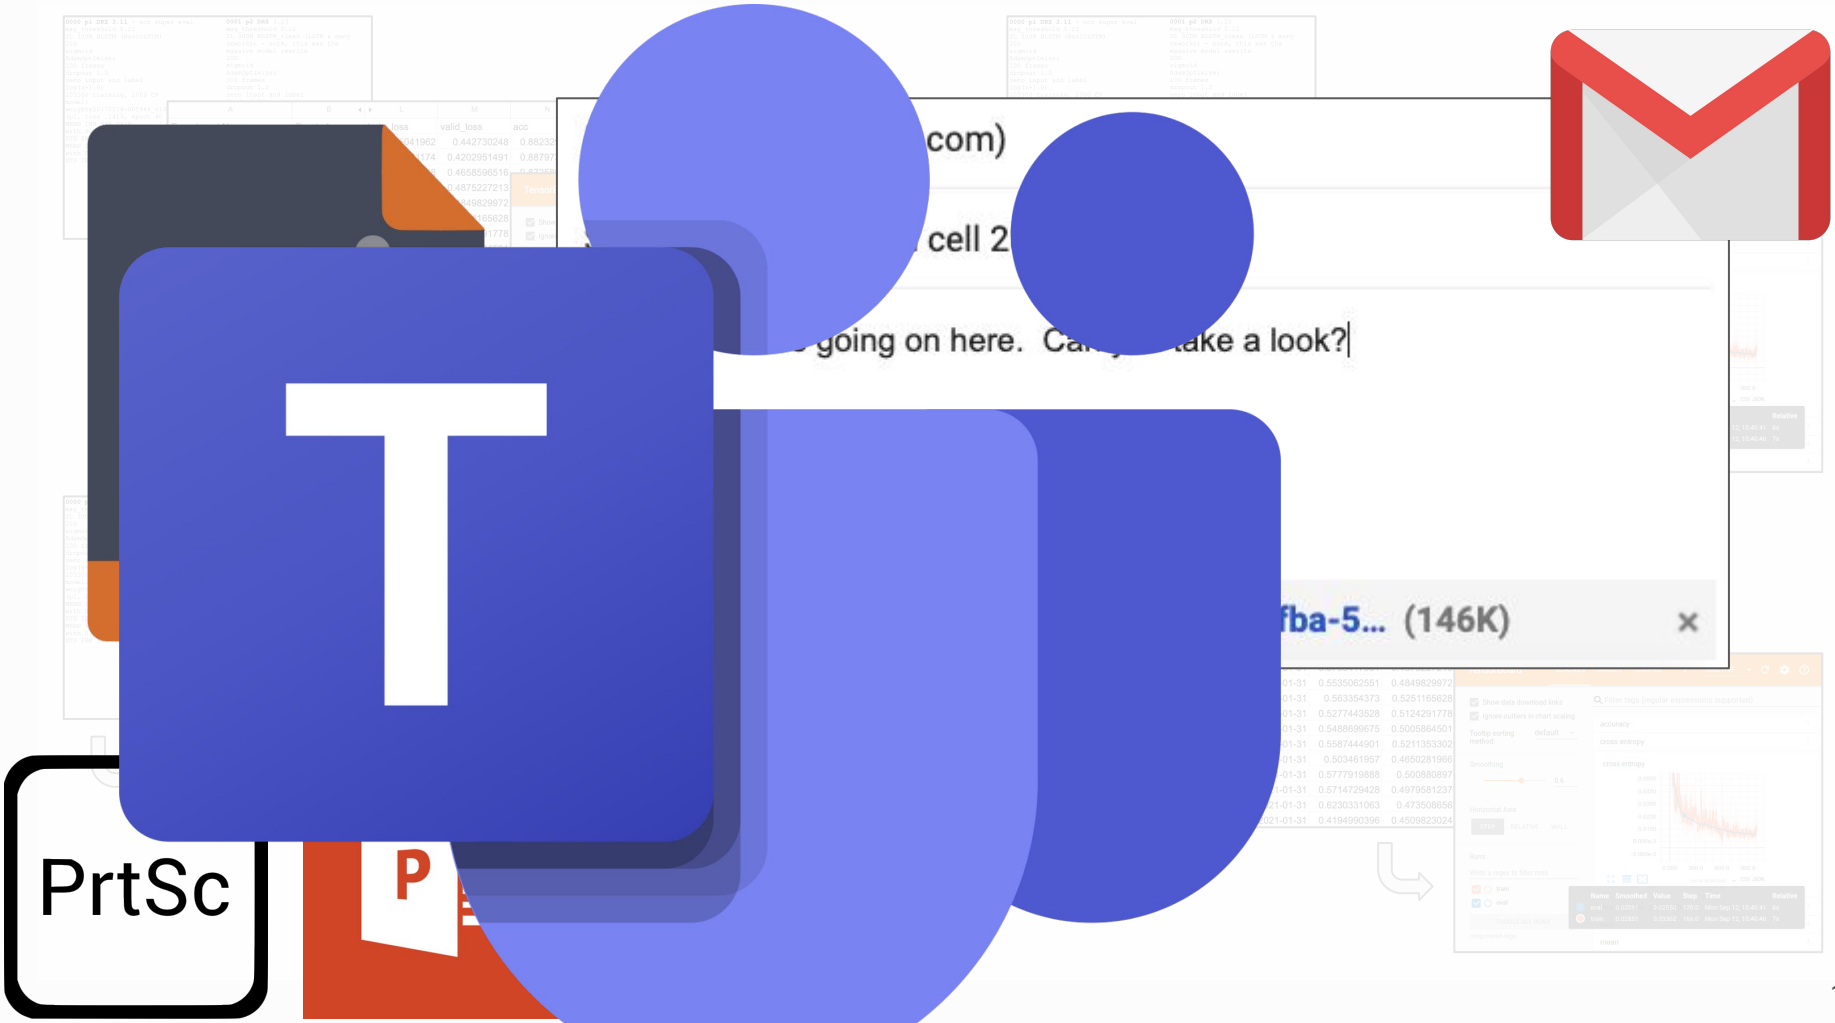
com)
cell 2
going on here. Ca take a look?
fba-5... (146K)
PrtSc


PrtSc

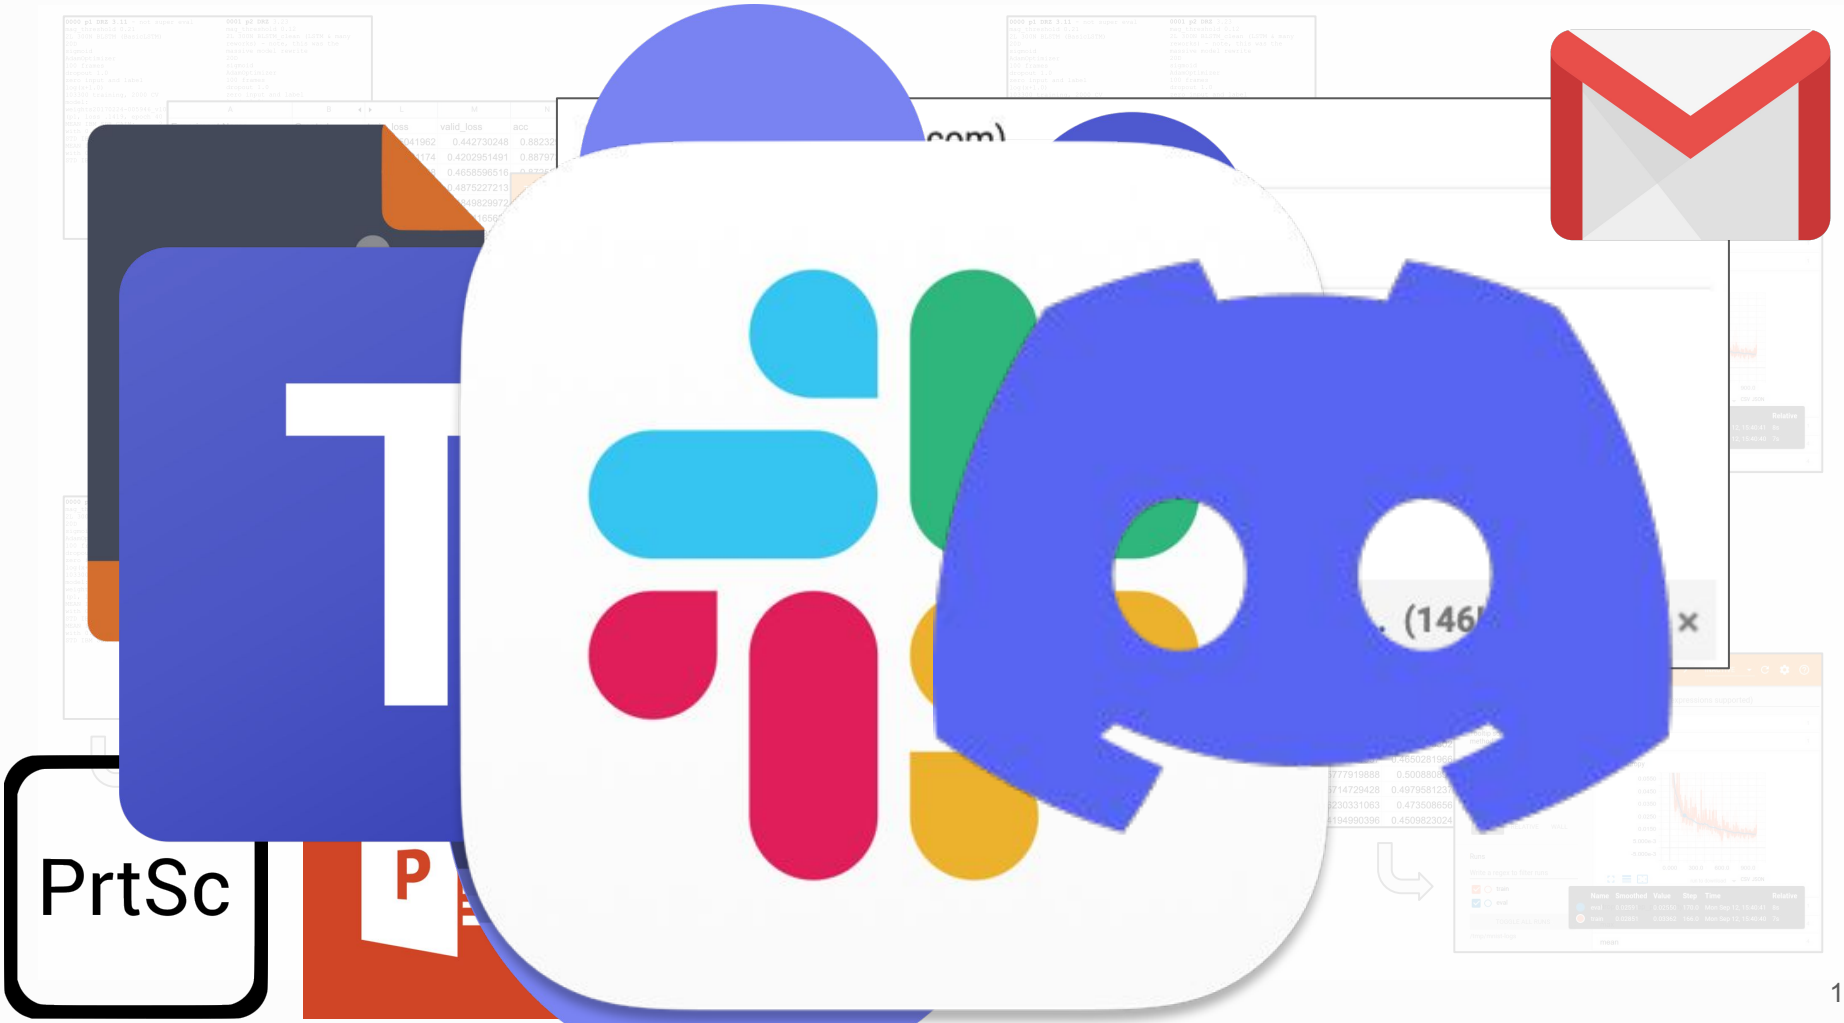
PrtSc


THIS IS FINE.

Weights & Biases

Just Stop

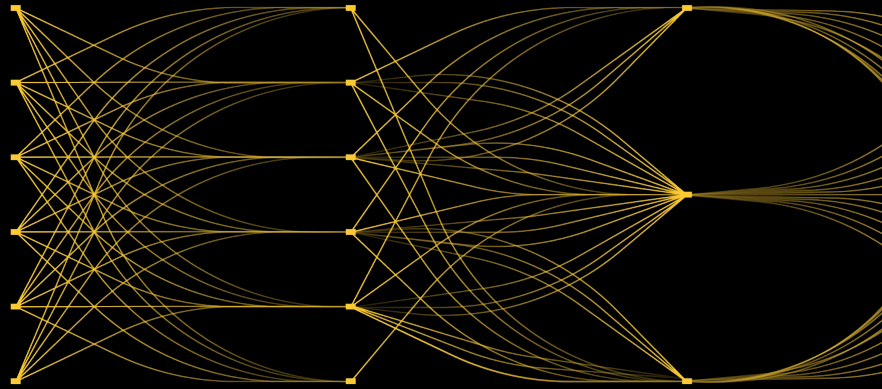

October 2022

Three principles of an ideal ML workflow

Rapidly iterate

to continuously refine
and optimize models

Reproduce

to reduce key-person
dependencies

Collaborate

to ensure knowledge transfer
across the organization

Three principles of an ideal ML workflow

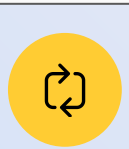

Rapidly iterate

to continuously refine and optimize models

Reproduce

to reduce key-person dependencies

Collaborate

to ensure knowledge transfer across the organization

A system of record for all ML workflows

A system of record for all ML workflows

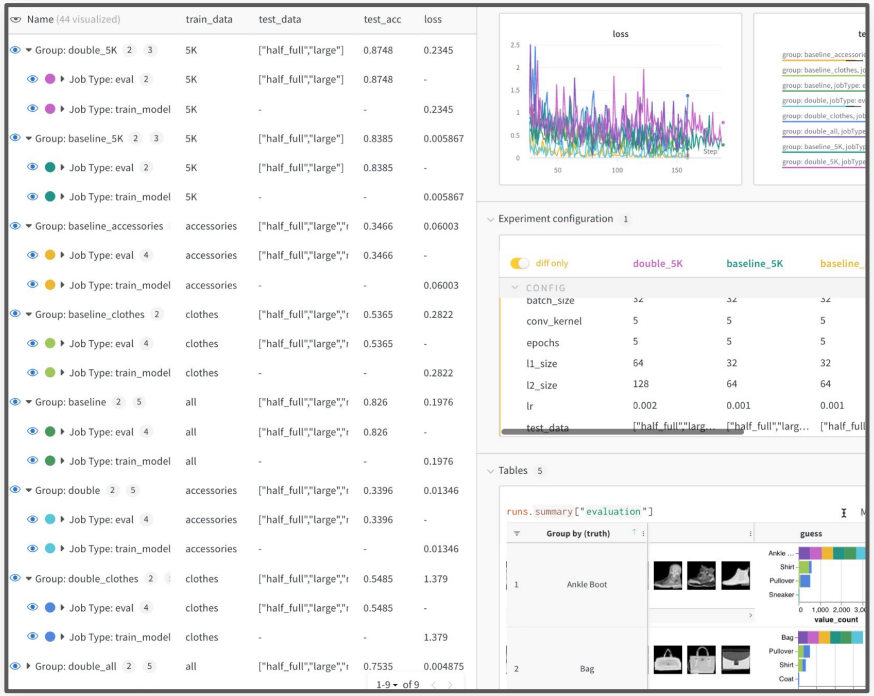

A system of record for all ML workflows

Get started in 60 seconds

```
!pip install wandb    # Install W&B

wandb.init()          # Start experiment
wandb.log(metrics)    # Log metrics + more!
```

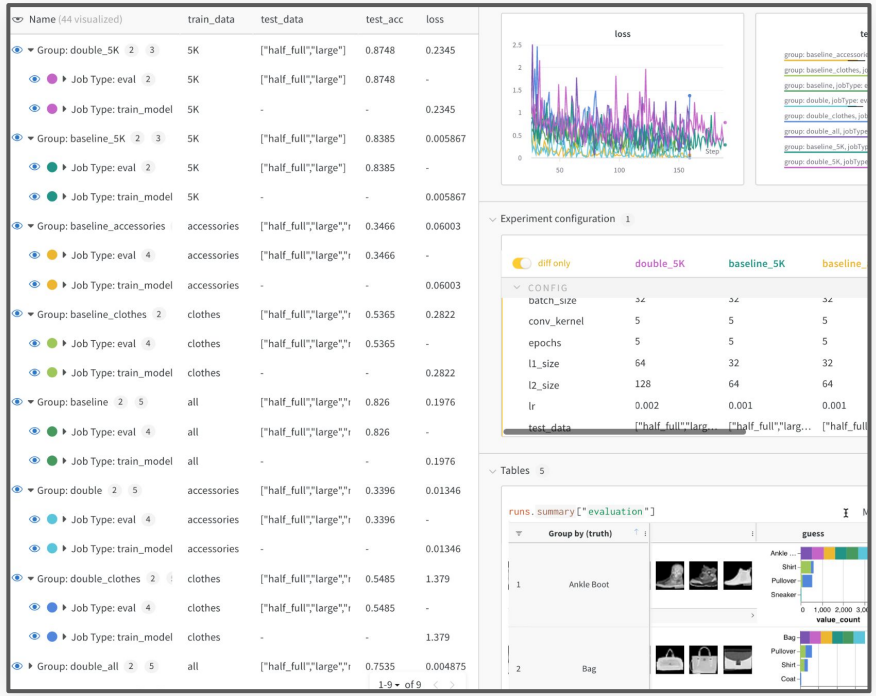

Yes, you really can get started in 60 seconds.

A system of record for `wandb.Image`

A system of record for `wandb.Object3D`

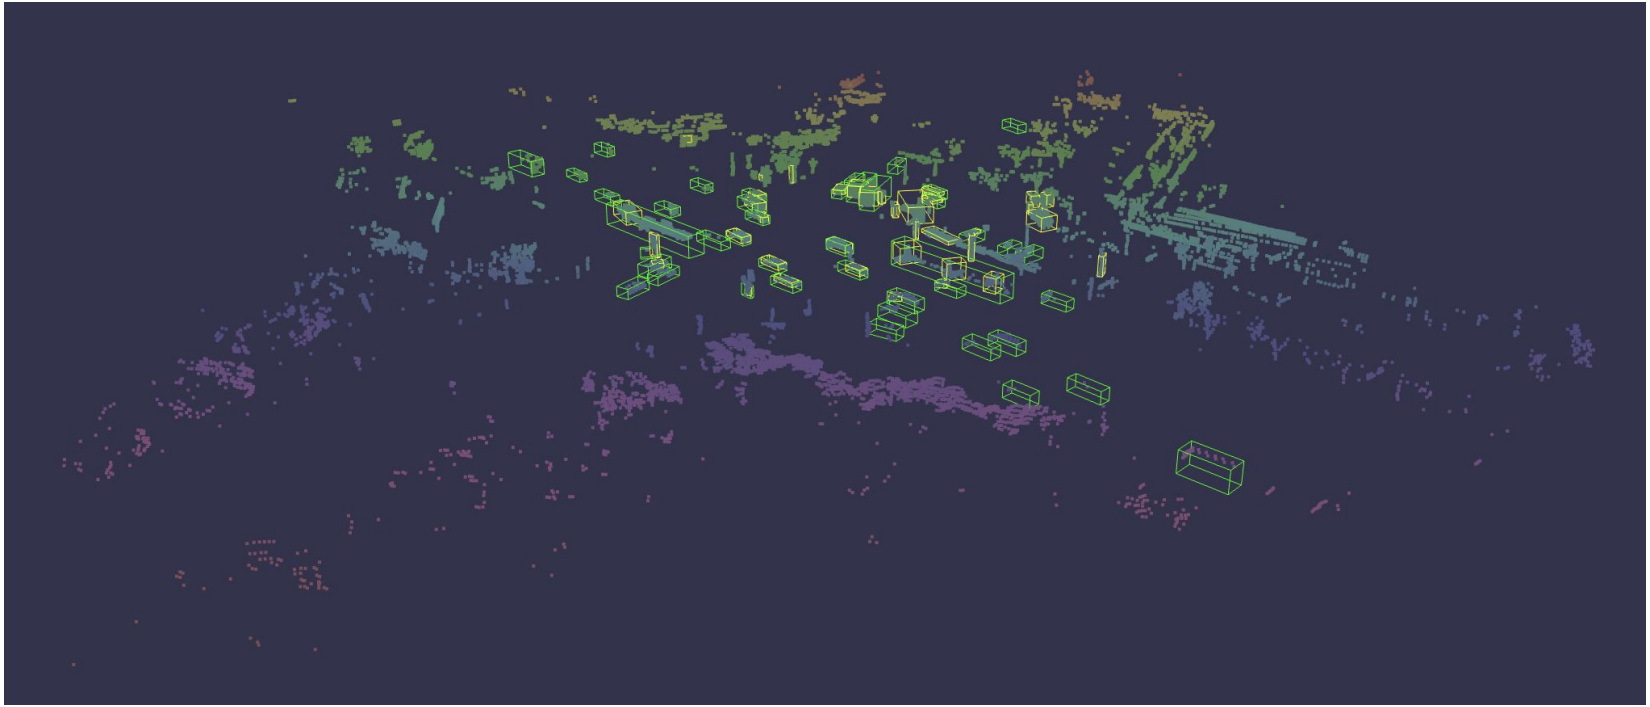

A system of record for `wandb.Audio`

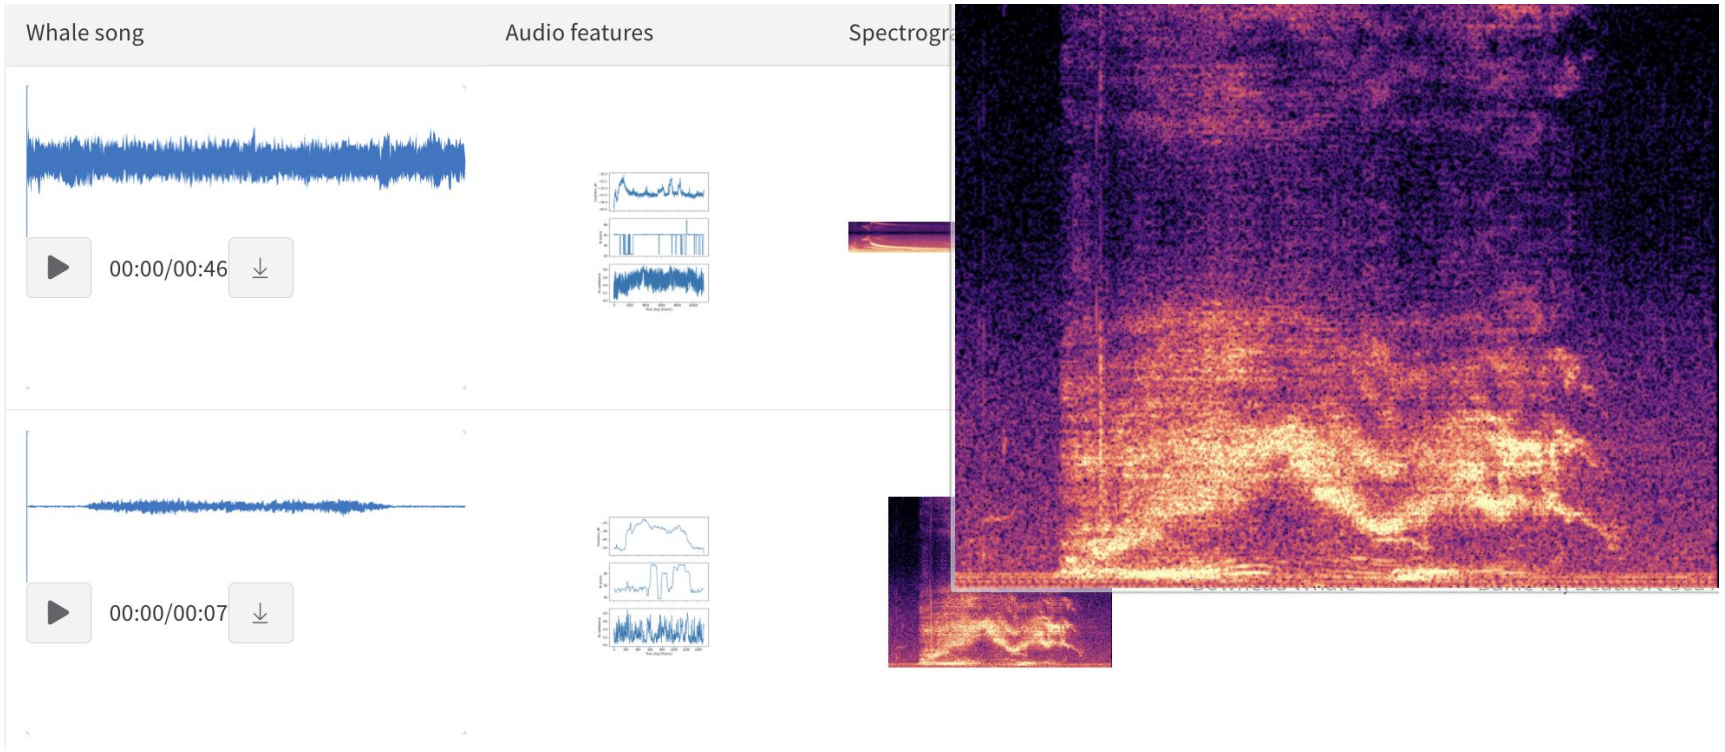

A system of record for `wandb.Video`

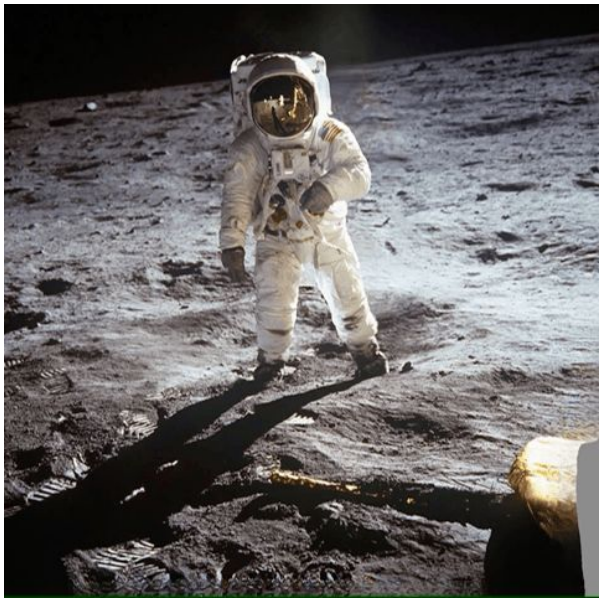

A system of record for `wandb.plots.POS`

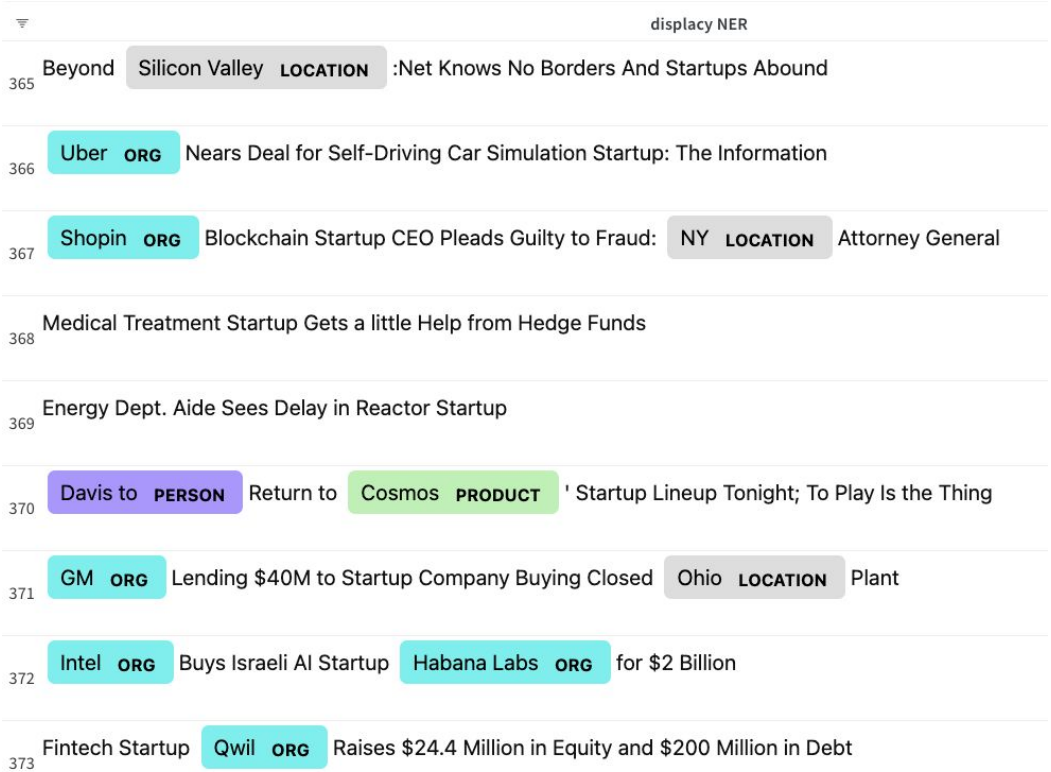

A system of record for `wandb.Html`

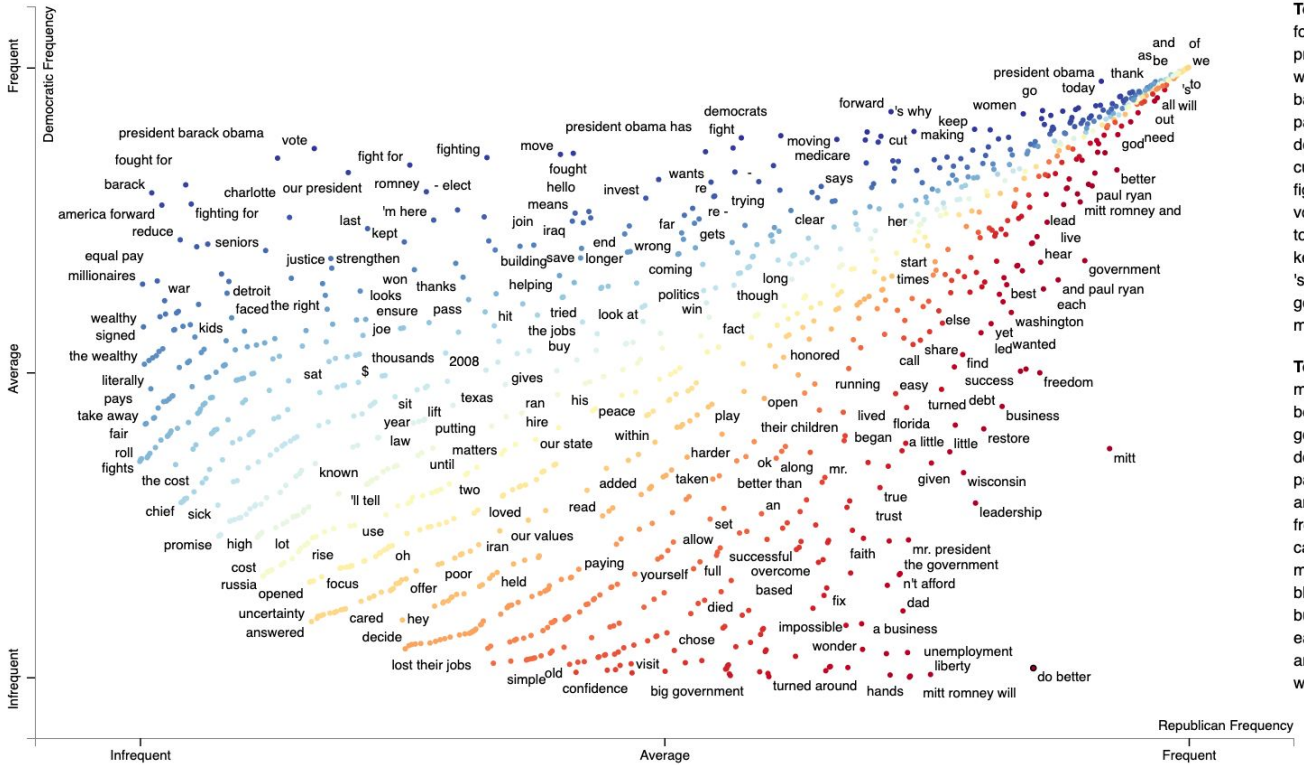

Top Democratic	Characteristic
forward	obamacare
president obama	mitt
women	barack
barack obama	obama
pay	romney
democrats	billionaires
cut	millionaires
fight	bless
vote	risked
together	gentlemen
keep	understands
's why	wo
go	fought
medicare	republicans
	elect
Top Republican	bankrupt
mitt	prosper
better	tonight
government	repeal
do better	democrats
paul ryan	medicare
and paul ryan	proud
freedom	honored
came	thank
mitt romney and	struggled
bless	prosperity
built	deserve
each	believes
around	nominate
wanted	pledged

A system of record for all types of data!

```
wandb.Image
wandb.Object3D
wandb.log( wandb.Molecule )
wandb.Video
wandb.Html
```

COVID-19 main protease in complex with Z31792168

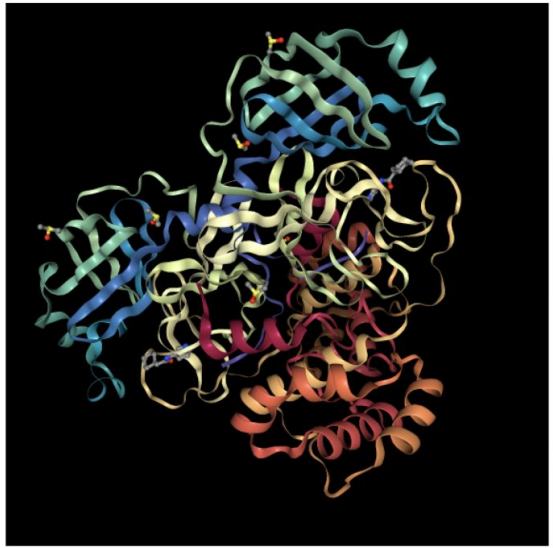

For a complete set of types, visit docs.wandb.ai

Ask questions about your data and models

Flask

web development,
one drop at a time

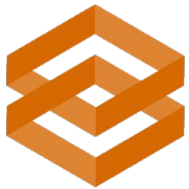

gradio

Streamlit

plotly | Dash

Ask questions about your data (and models)

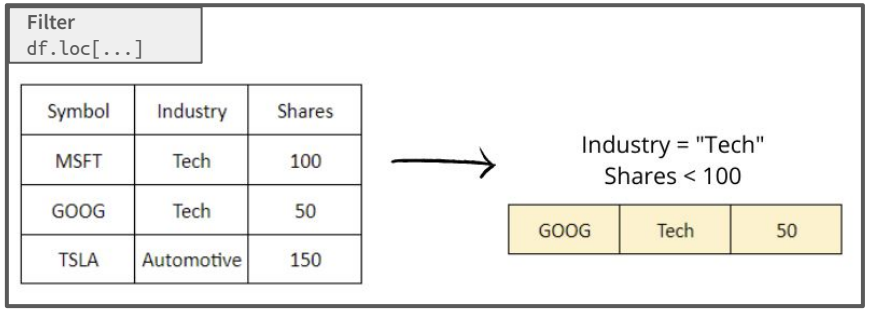

Ask questions about your data (and models)

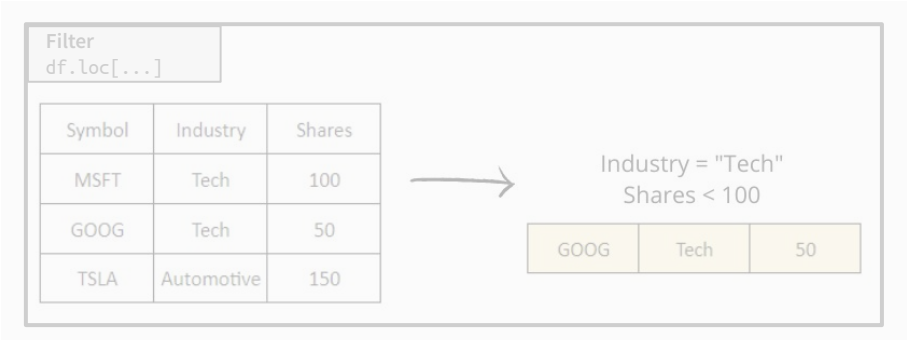

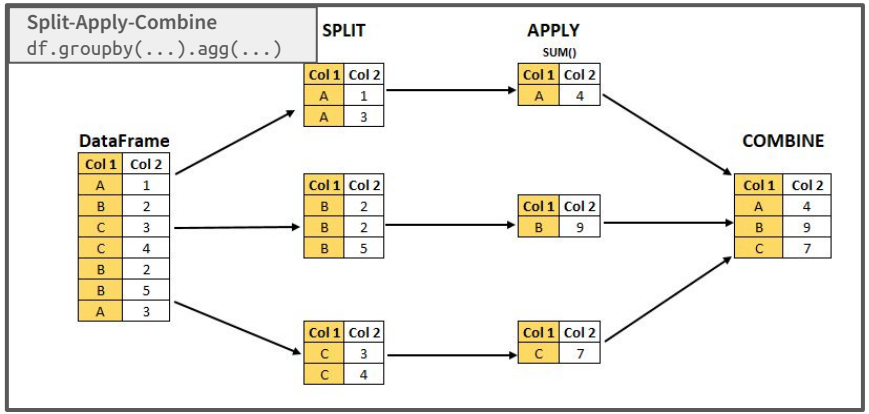

Ask questions about your data (and models)

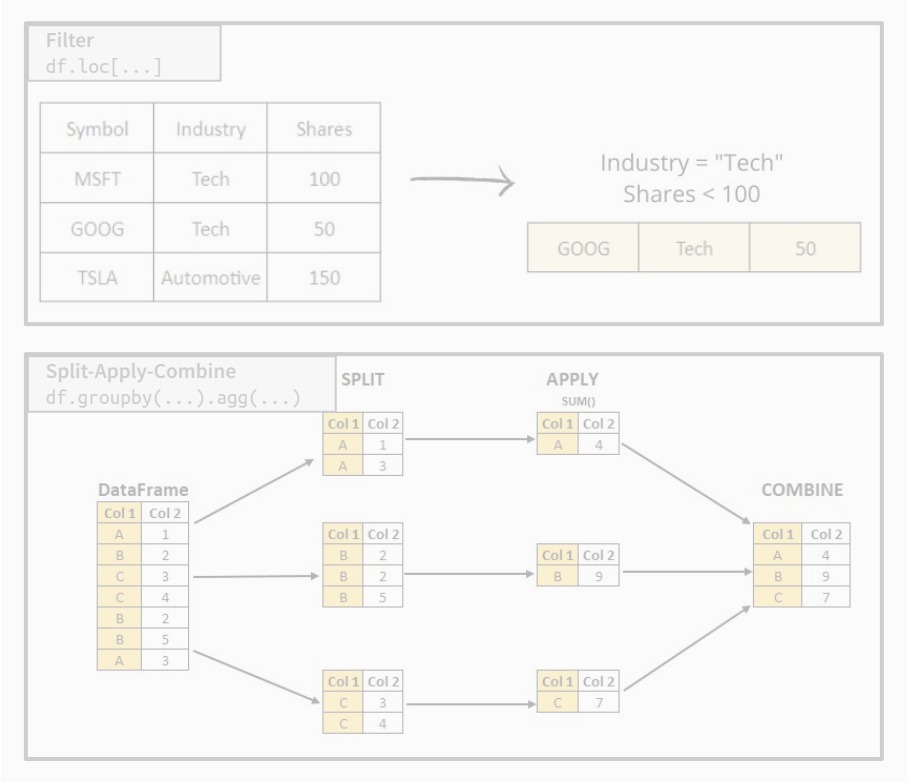

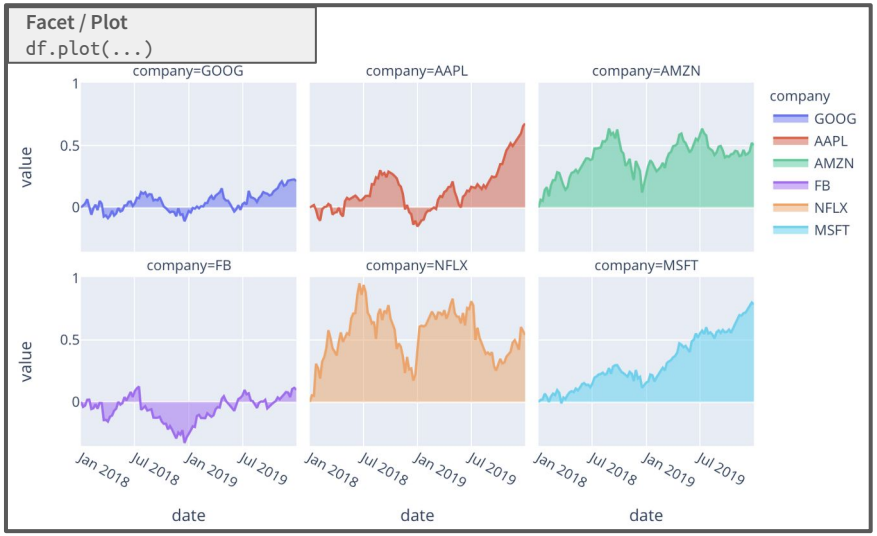

How do I ask questions about non-tabular data?

Ask questions about your data (and models)

❌ Classic DataFrames

	img	label
0	image mode=RGB size=500x333 at 0x7F56F6360ED0>	samoyed
1	image mode=RGB size=443x435 at 0x7F56F6379BD0>	shiba_inu
2	image mode=RGB size=600x437 at 0x7F56F6379E50>	Egyptian_Mau
3	image mode=RGB size=375x500 at 0x7F56F6379F90>	Birman
4	image mode=RGB size=500x375 at 0x7F56F637B350>	great_pyrenees
...	...	...
7344	image mode=RGB size=500x335 at 0x7F56F623BD50>	Bengal
7345	image mode=RGB size=403x500 at 0x7F56F623BF90>	leonberger
7346	image mode=RGB size=500x375 at 0x7F56F6241390>	beagle
7347	image mode=RGB size=192x288 at 0x7F56F62416D0>	Abyssinian
7348	image mode=RGB size=288x300 at 0x7F56F6241910>	american_pit_bull_terrier

Ask questions about your data (and models)

❌ Classic DataFrames

	img	label
	image mode=RGB size=500x333 at 0x7F56F6360ED0>	samoyed
	image mode=RGB size=443x435 at 0x7F56F6379BD0>	shiba_inu
	image mode=RGB size=600x437 at 0x7F56F6379E50>	Egyptian_Mau
	image mode=RGB size=375x500 at 0x7F56F6379F90>	Birman
	image mode=RGB size=500x375 at 0x7F56F637B350>	great_pyrenees
	...	...
	image mode=RGB size=500x335 at 0x7F56F623BD50>	Bengal
	image mode=RGB size=403x500 at 0x7F56F623BF90>	leonberger
	image mode=RGB size=500x375 at 0x7F56F6241390>	beagle
	image mode=RGB size=192x288 at 0x7F56F62416D0>	Abyssinian
	image mode=RGB size=288x300 at 0x7F56F6241910>	american_pit_bull_terrier

✅ wandb.Table

	img	label
1		samoyed
2	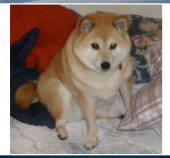	shiba_inu
3		Egyptian_Mau
4		Birman

Ask questions about your data (and models)

✅ wandb.Table

	img	label
1		samoyed
2	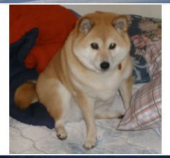	shiba_inu
3		Egyptian_Mau
4		Birman

`Table.groupby("label")`

	Group by (label)	img
1	japanese_chin	 6-10 of 24
2	Ragdoll	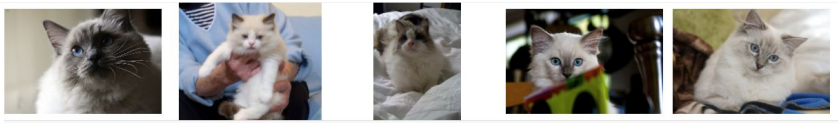 1-5 of 34
3	leonberger	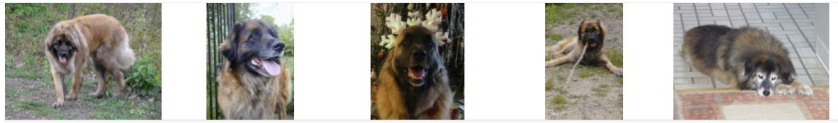 1-5 of 30
4	shiba_inu	 1-5 of 25

How do I ask questions about training?

Sweeps

Easily and systematically search hyperparameters

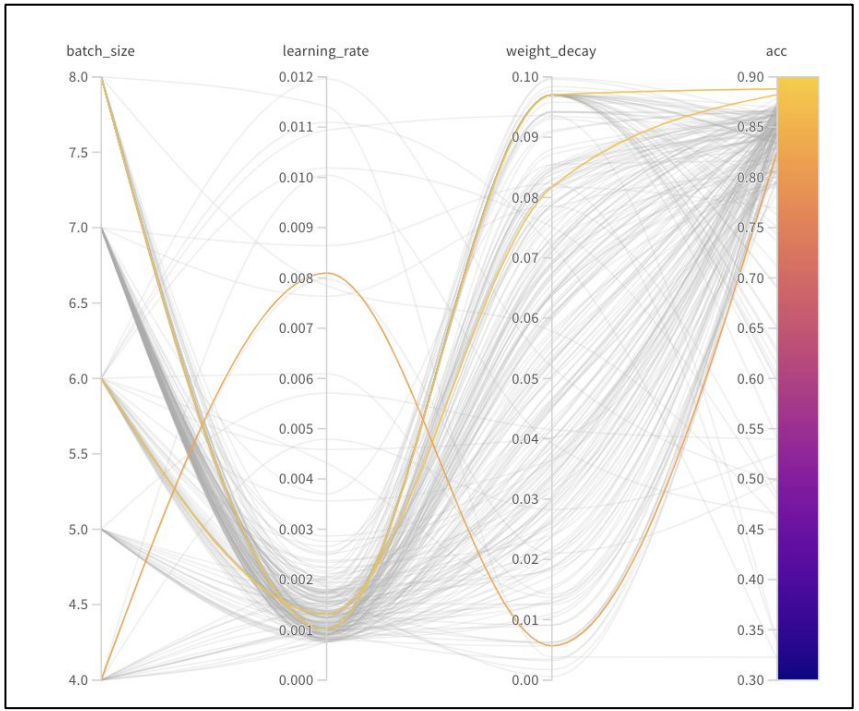

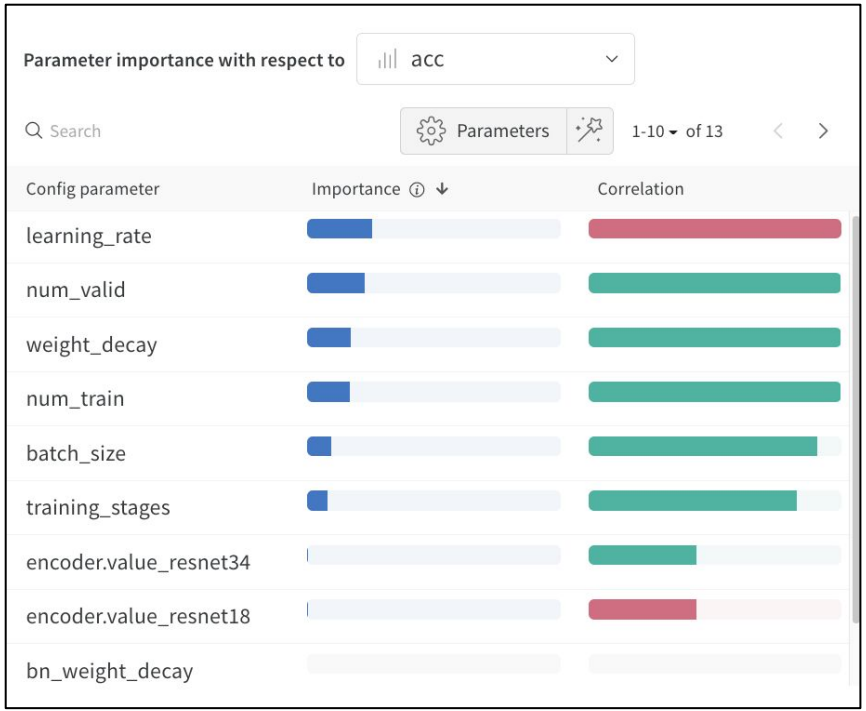

Ask questions about your data (and models)

```
program: train.py
method: bayes
metric:
 name: valid_loss
 goal: minimize
parameters:
 batch_size:
   values: [32, 64]
 mixup_alpha:
   values: [0.2, 0.5, 0.8]
 optimizer:
   values: ["adam", "ranger"]
 encoder:
   values: ["resnet18", "resnet34", "resnet50", "resnet101"]
```

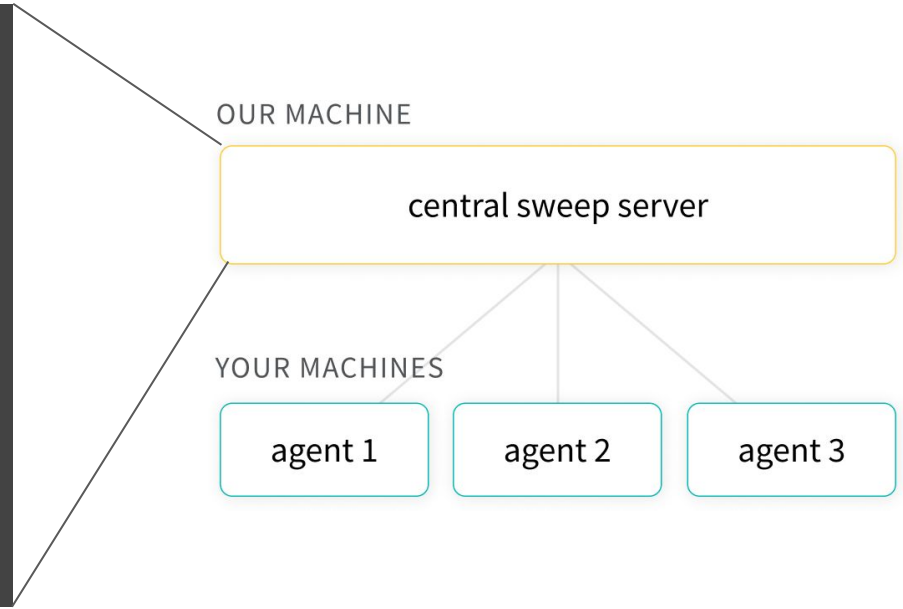

Three principles of an ideal ML workflow

Rapidly iterate

to continuously refine and optimize models

Reproduce

to reduce key-person dependencies

Collaborate

to ensure knowledge transfer across the organization

Artifacts – Version data and models across your pipeline

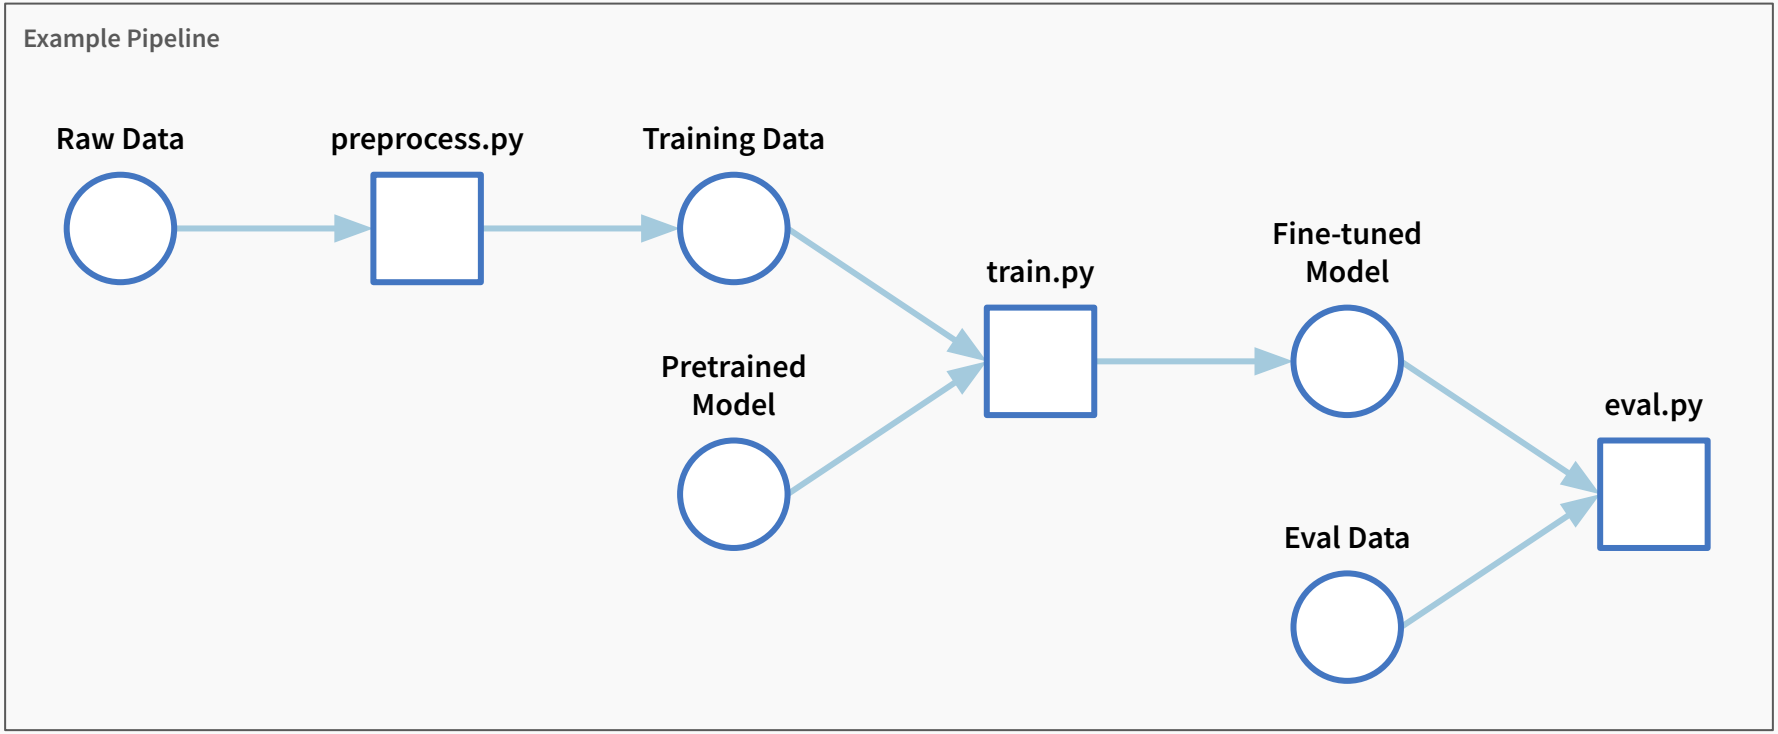

Bad predictions!

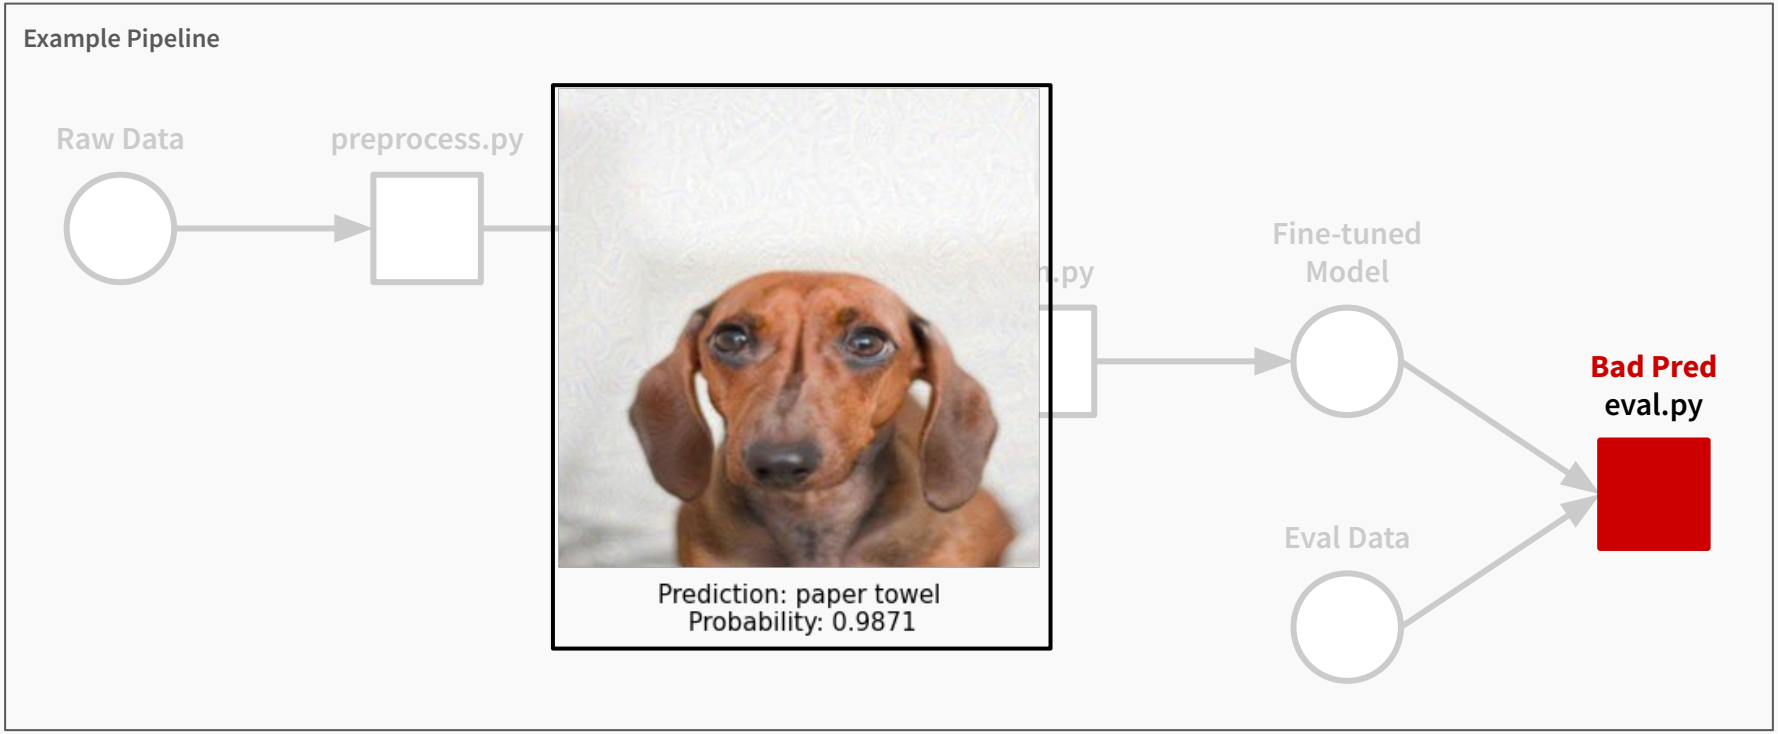

Bad predictions because of **training?**

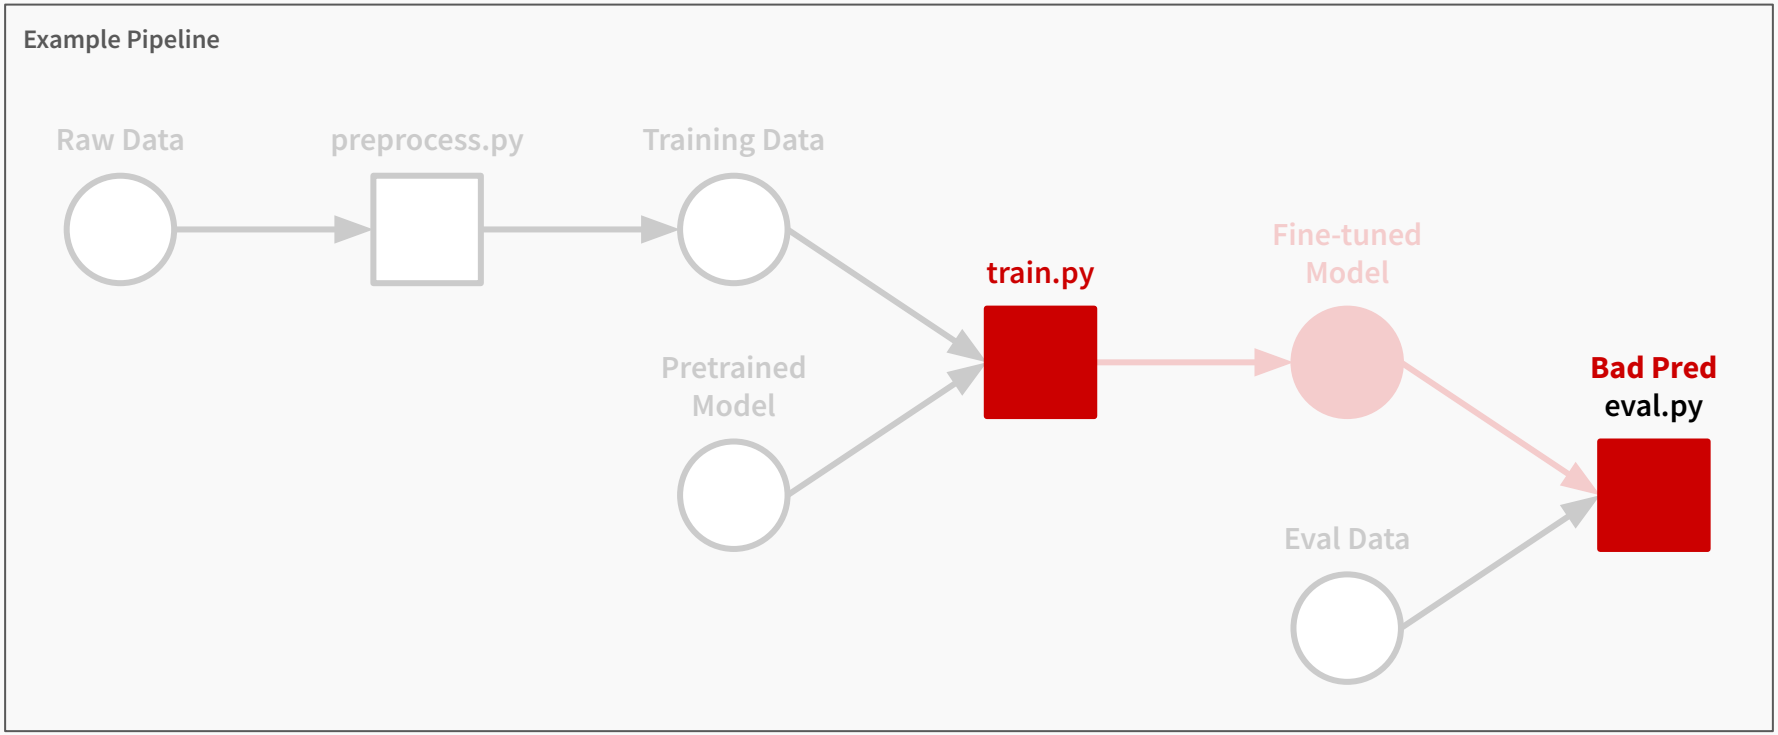

Bad predictions because of **pretrained models?**

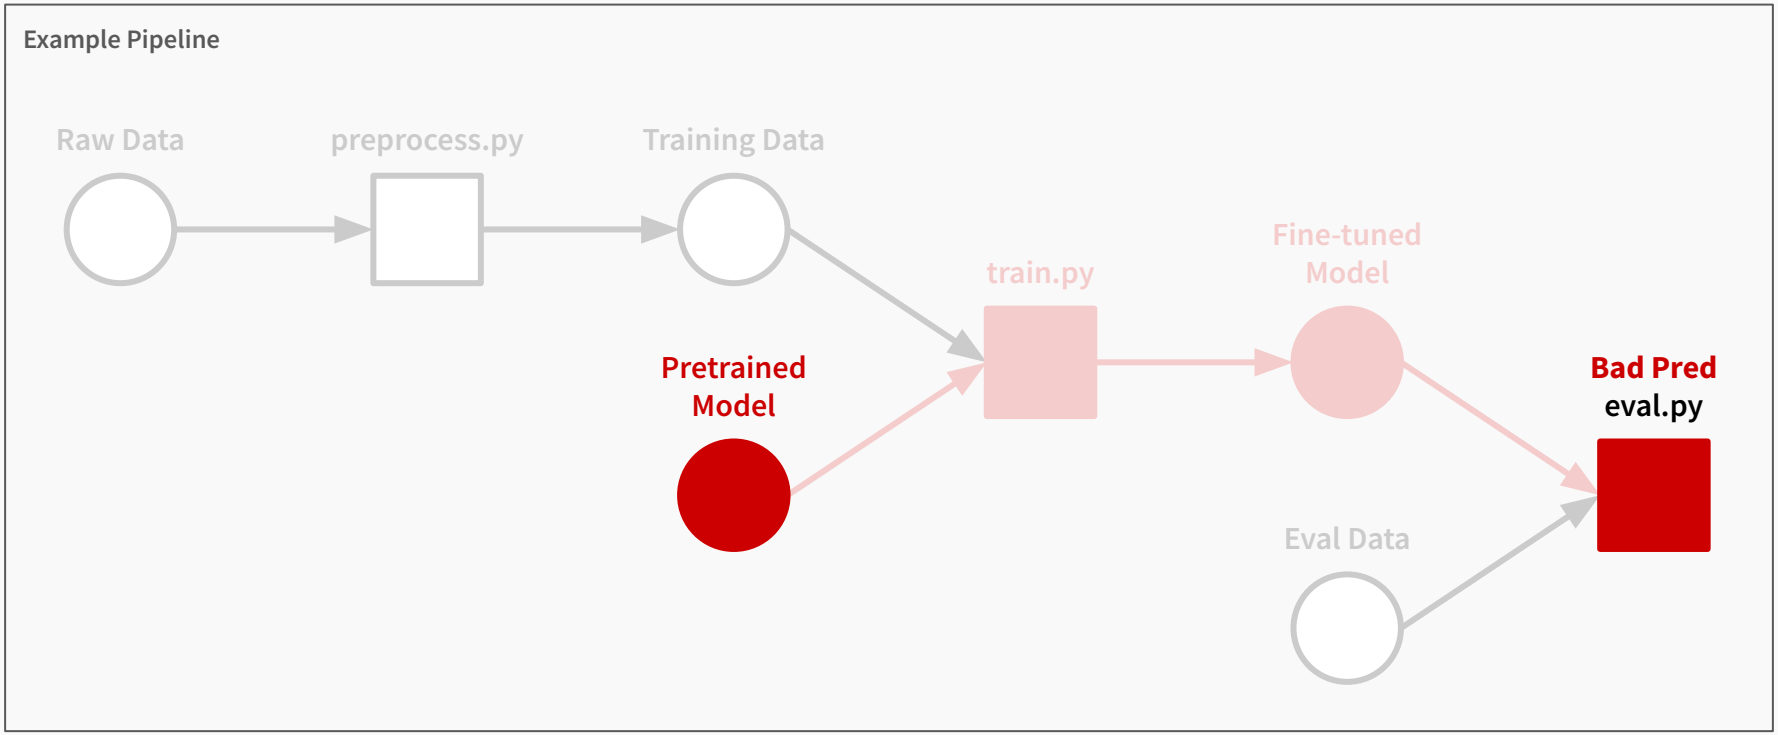

Bad predictions because of **preprocessing?**

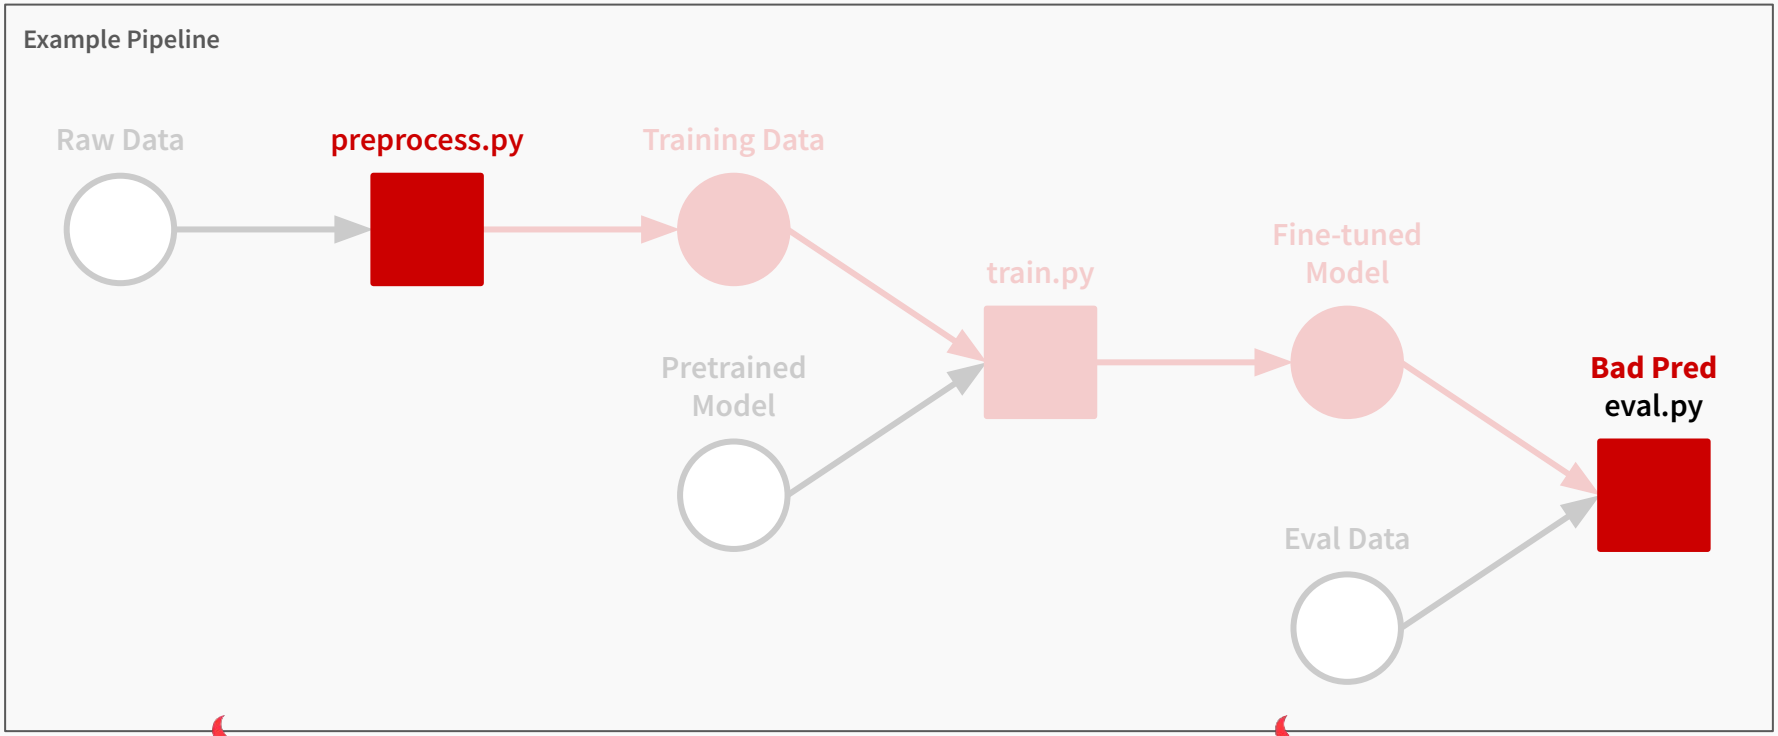

Bad predictions because of **raw data?**

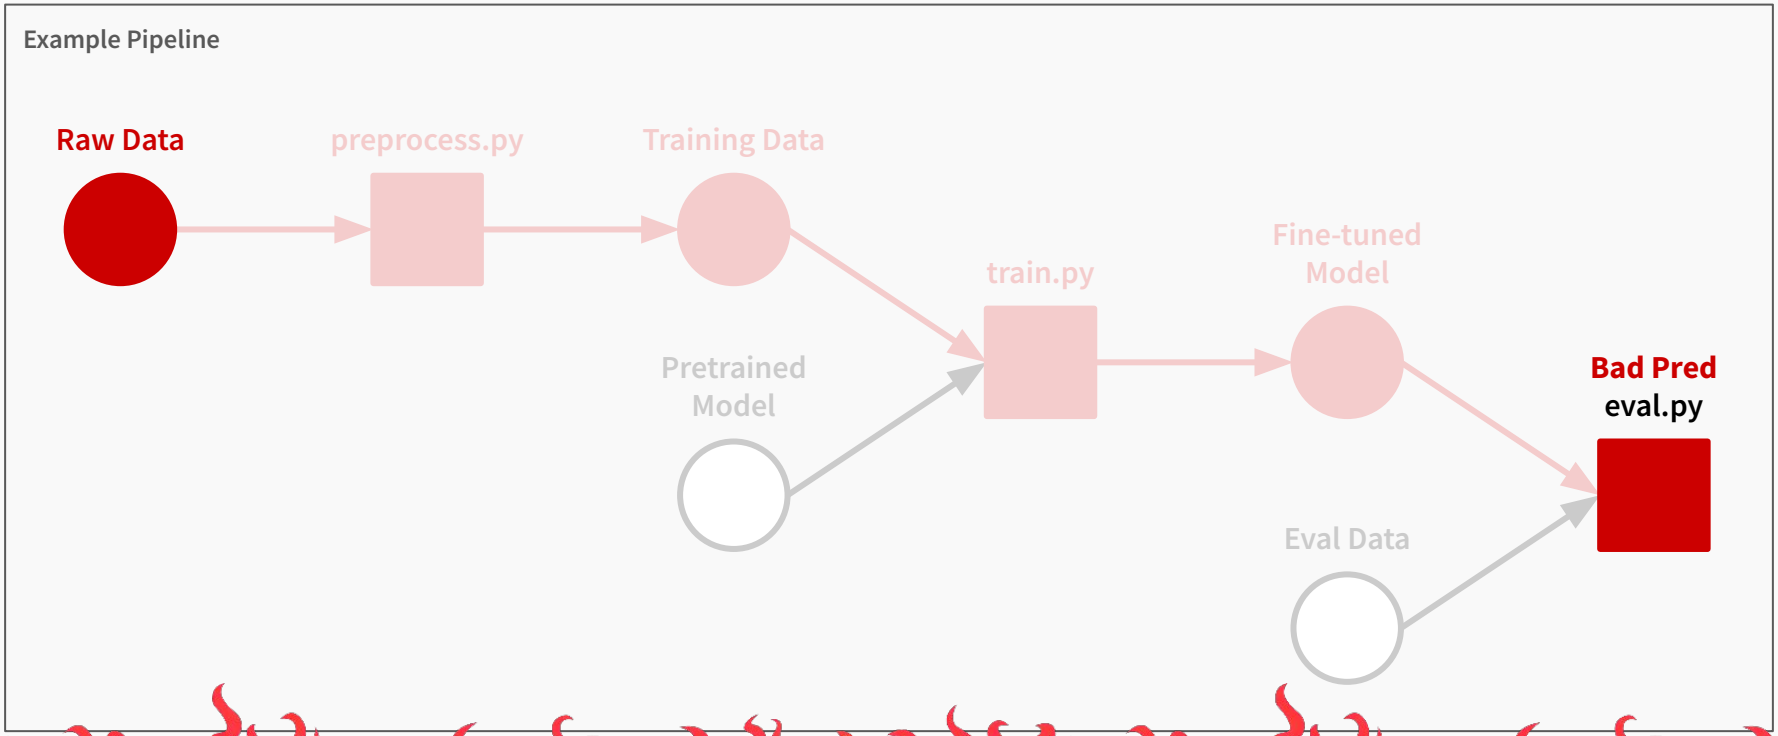

Artifacts = Version data and models across your pipeline

Bad predictions because of **multiprocessing?**

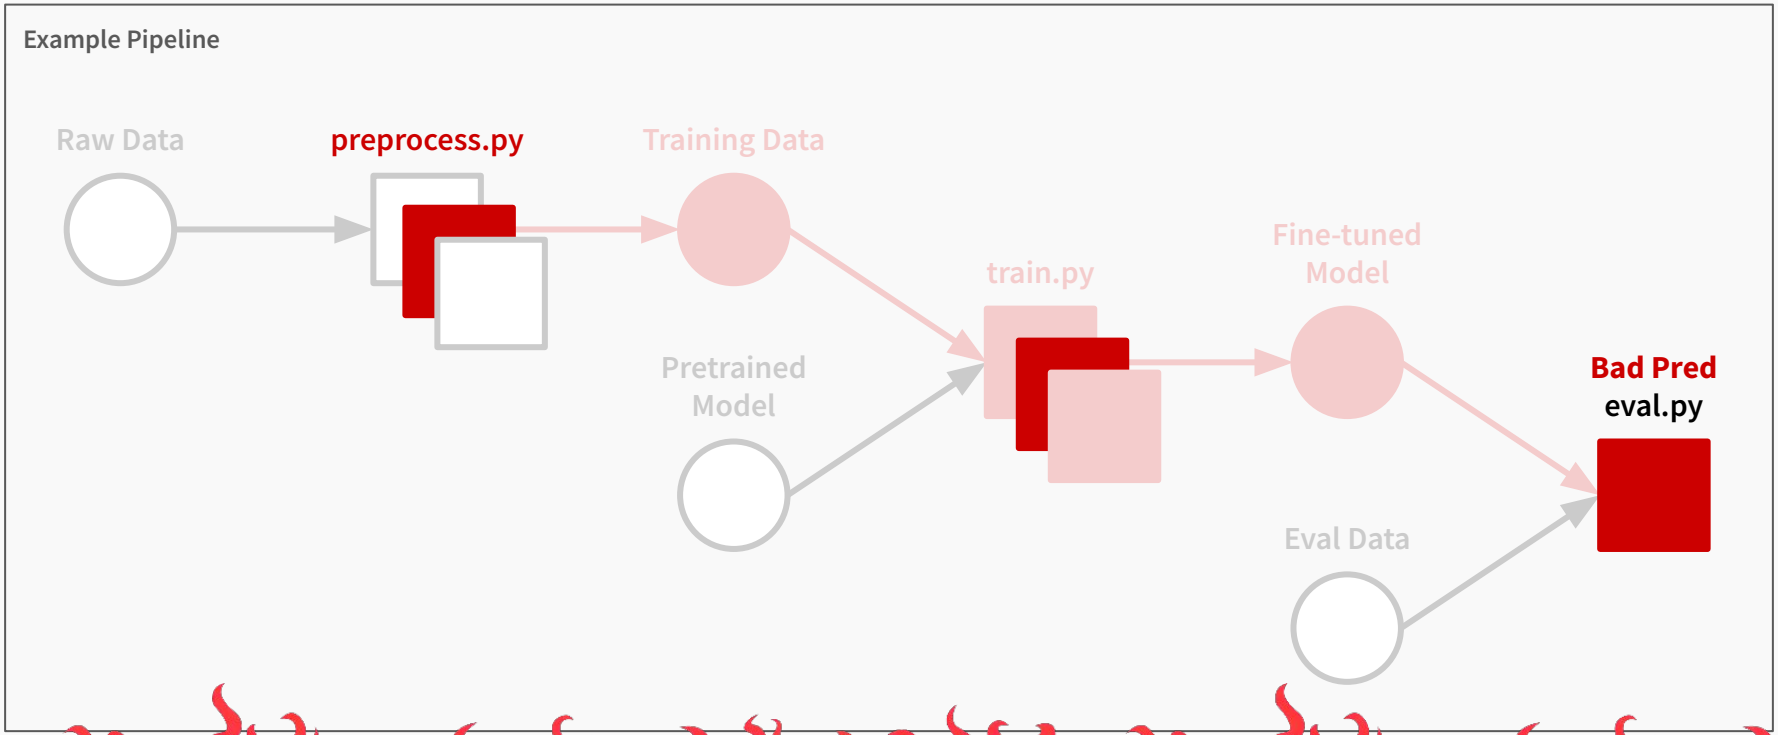


THIS IS FINE.

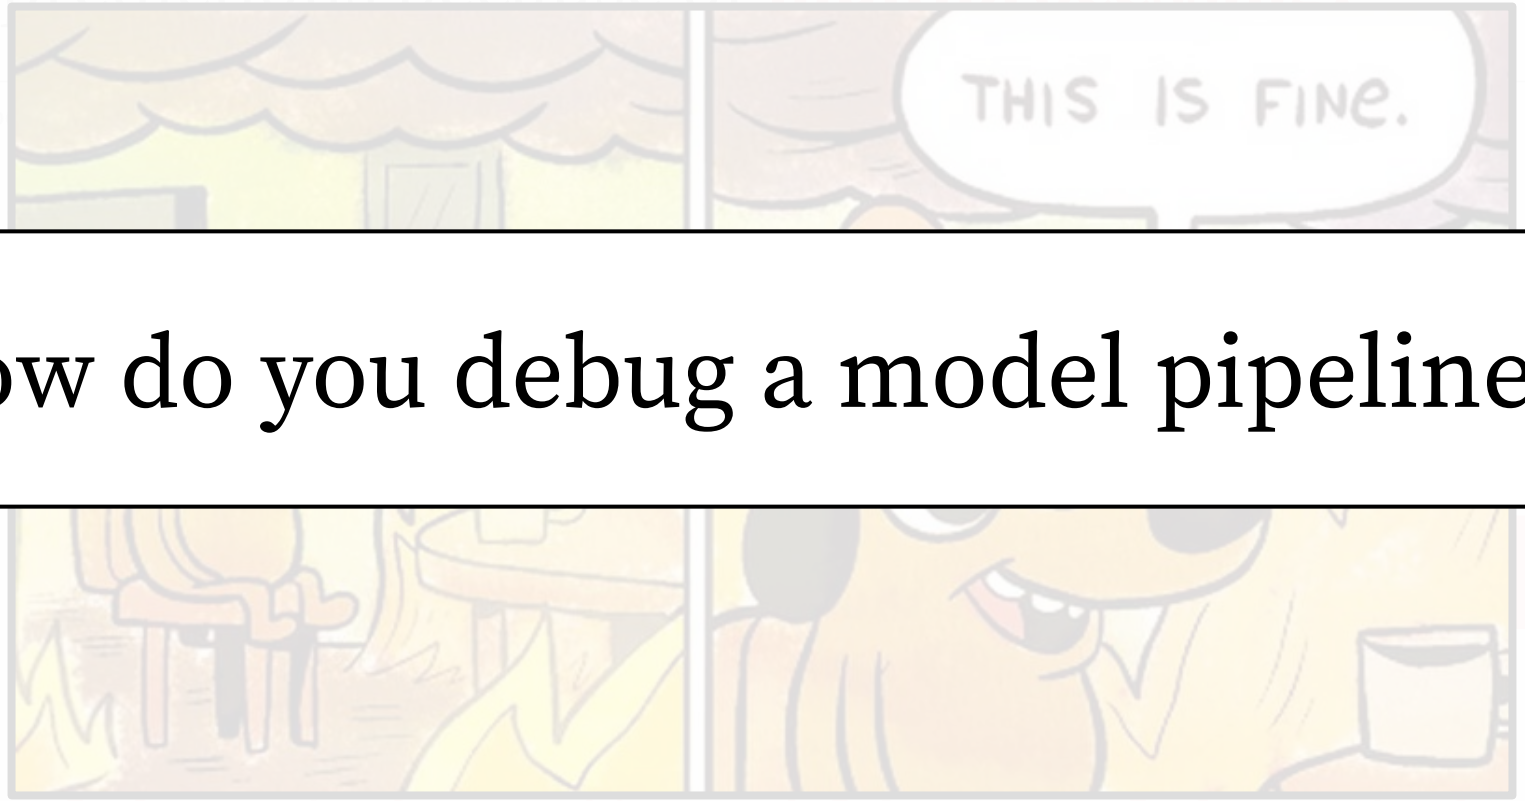

How do you debug a model pipeline?

Check the Code

Problem with code? More than `git diff`

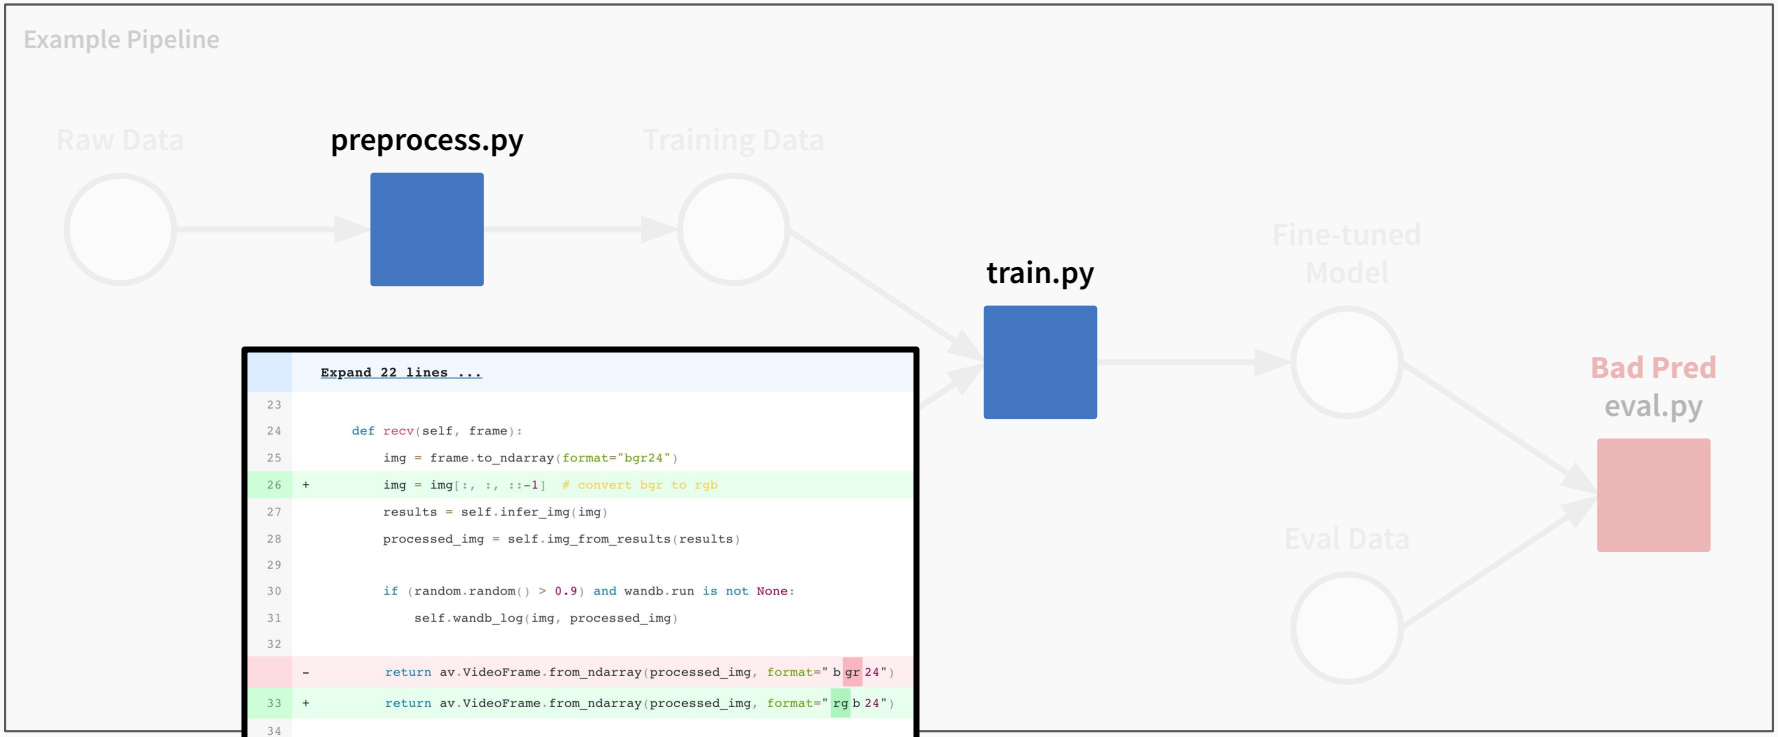


```
Expand 22 lines ...

23
24      def recv(self, frame):
25          img = frame.to_ndarray(format="bgr24")
26  +       img = img[:, :, ::-1]  # convert bgr to rgb
27          results = self.infer_img(img)
28          processed_img = self.img_from_results(results)
29
30          if (random.random() > 0.9) and wandb.run is not None:
31              self.wandb_log(img, processed_img)
32
    -       return av.VideoFrame.from_ndarray(processed_img, format="bgr24")
33  +       return av.VideoFrame.from_ndarray(processed_img, format="rgb24")
34
35      def infer_img(self, img):
36          results = model(img)
```

What about the input/output artifacts?

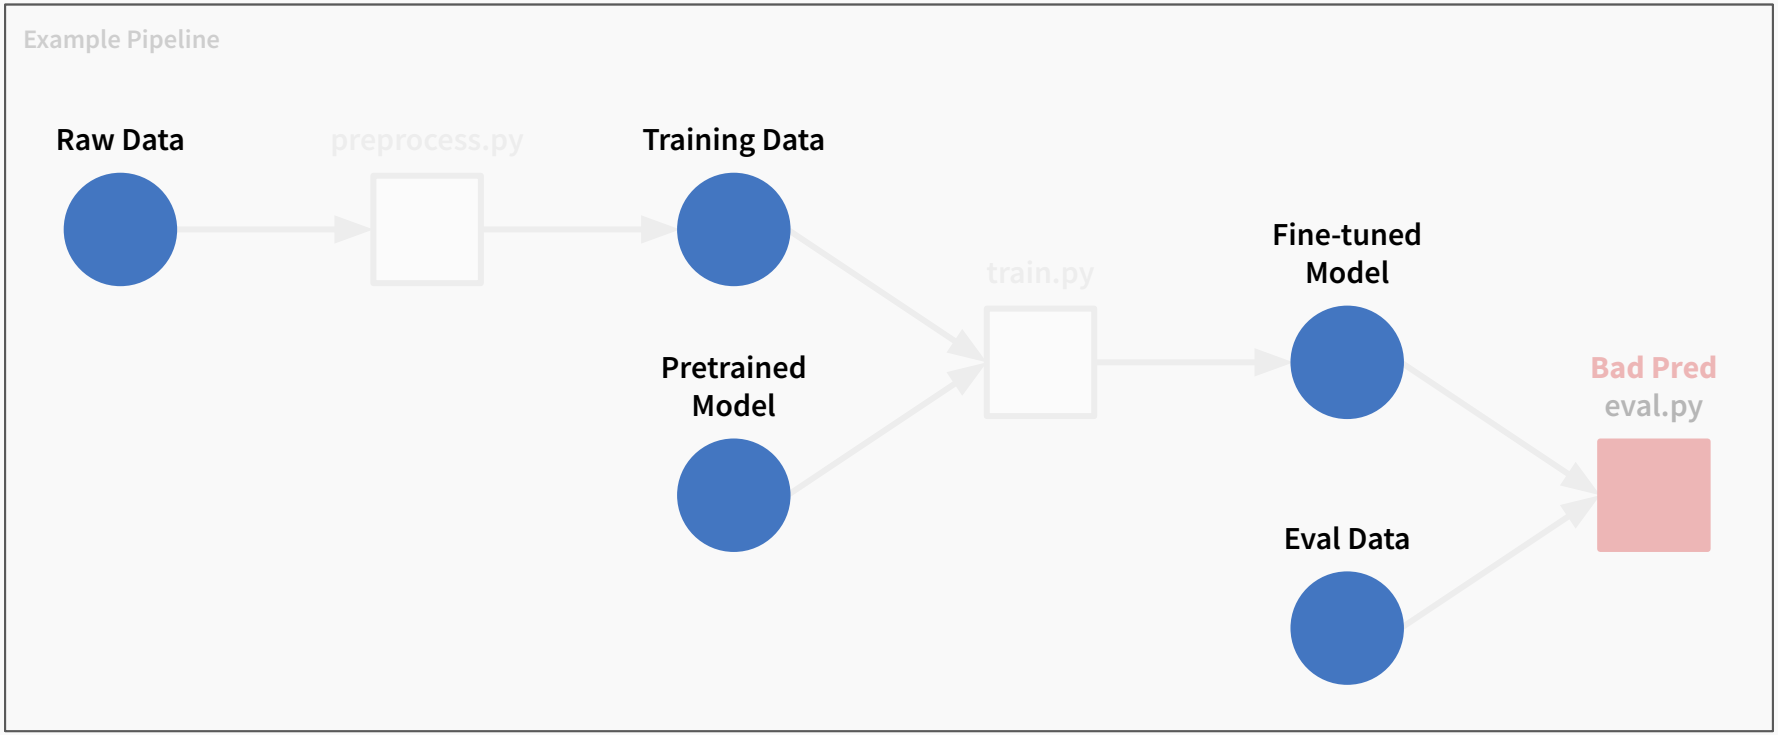

Check the inputs

Bad predictions because of **Raw Data**

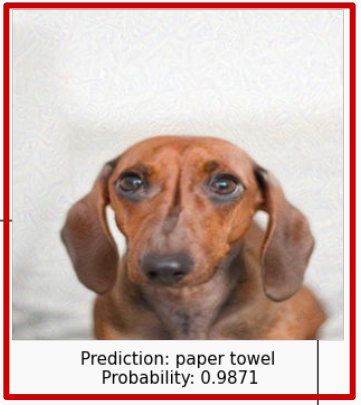

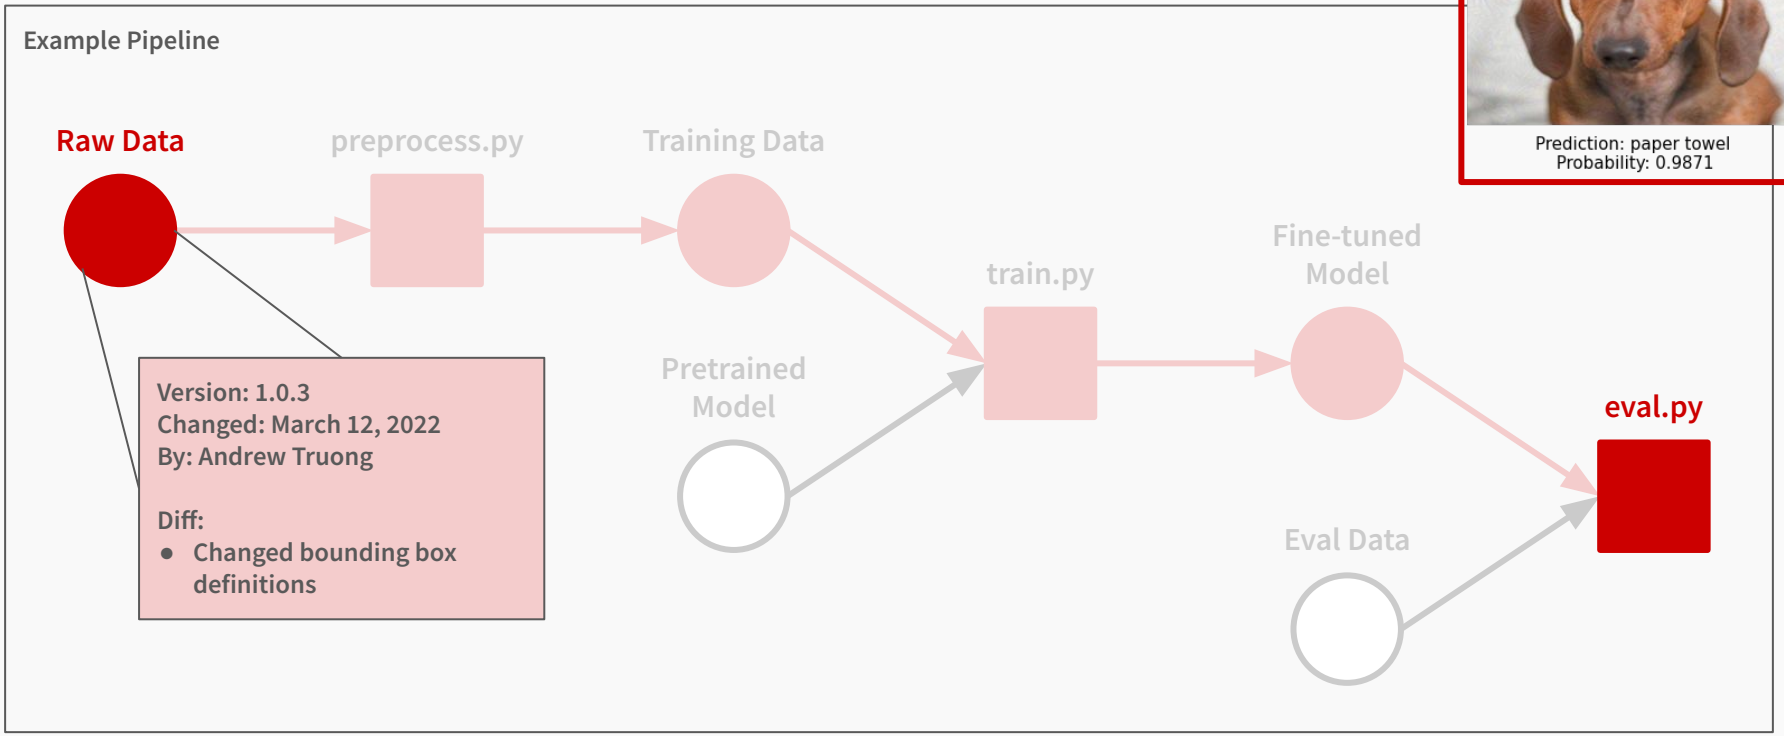

Good predictions after reverting Raw Data

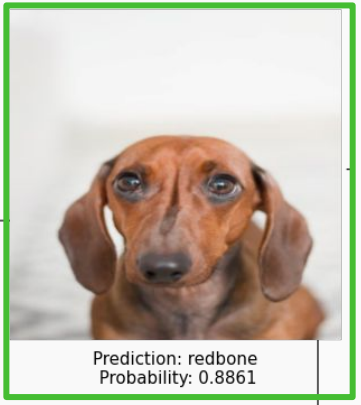

Example Pipeline

Raw Data

preprocess.py

Training Data

Pretrained Model

train.py

Fine-tuned Model

Eval Data

eval.py

Version: 1.0.2
Changed: March 11, 2022
By: Ben Sherman

Diff:
- Added 1713 new images
- Updated annotations.json

Check the outputs

Tables – DataFrames with rich media support

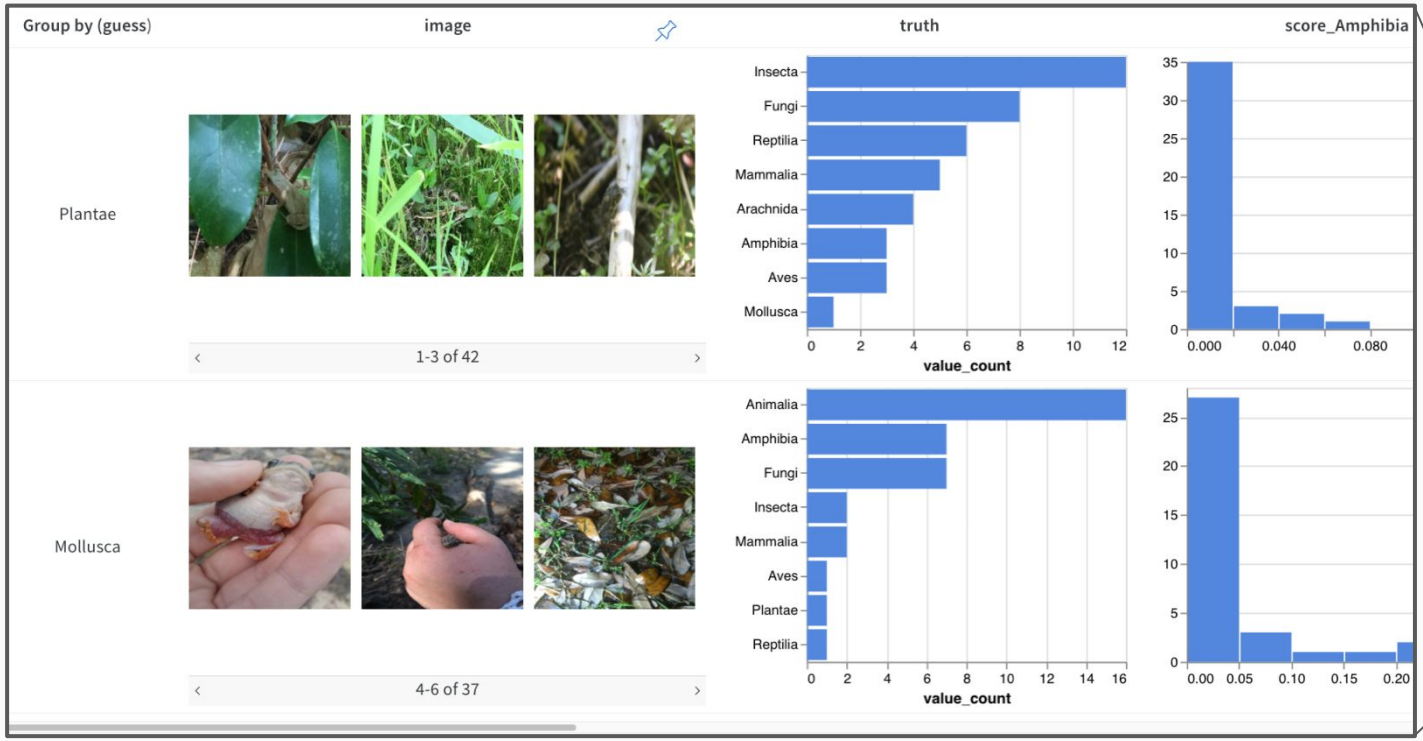

eval.py

Three principles of an ideal ML workflow

Rapidly iterate

to continuously refine and optimize models

Reproduce

to reduce key-person dependencies

Collaborate

to ensure knowledge transfer across the organization

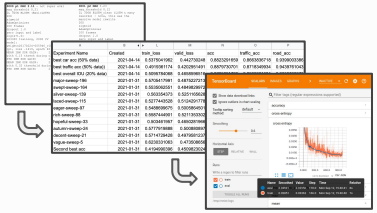
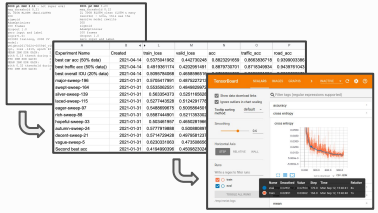
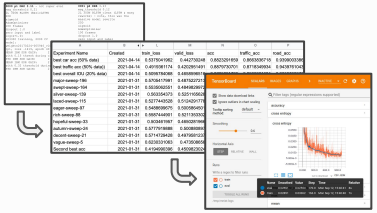
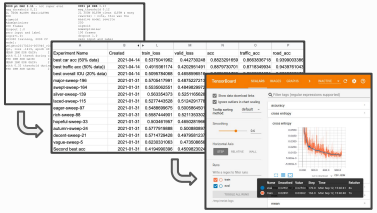
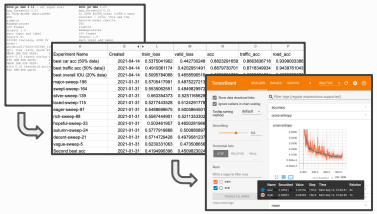
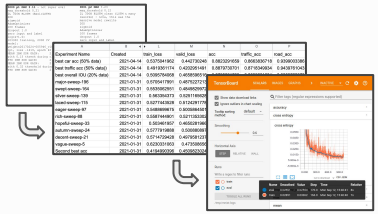
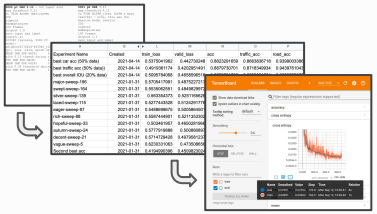
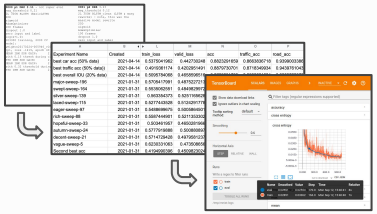
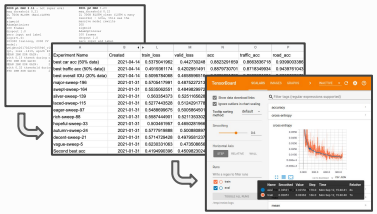
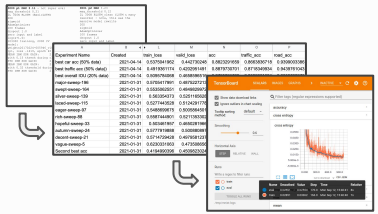
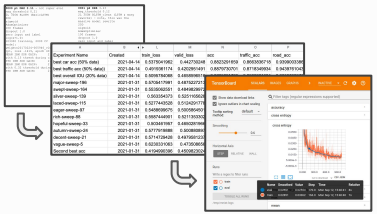
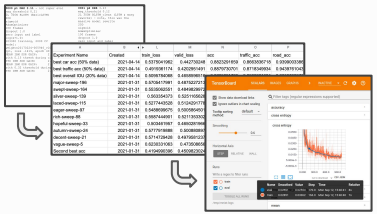
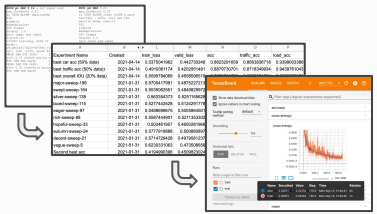
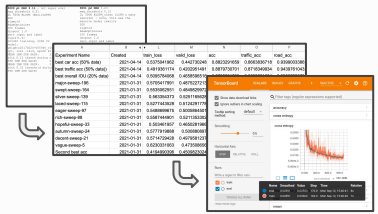
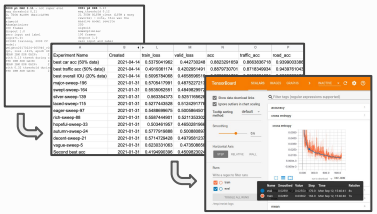
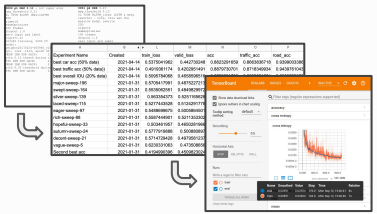

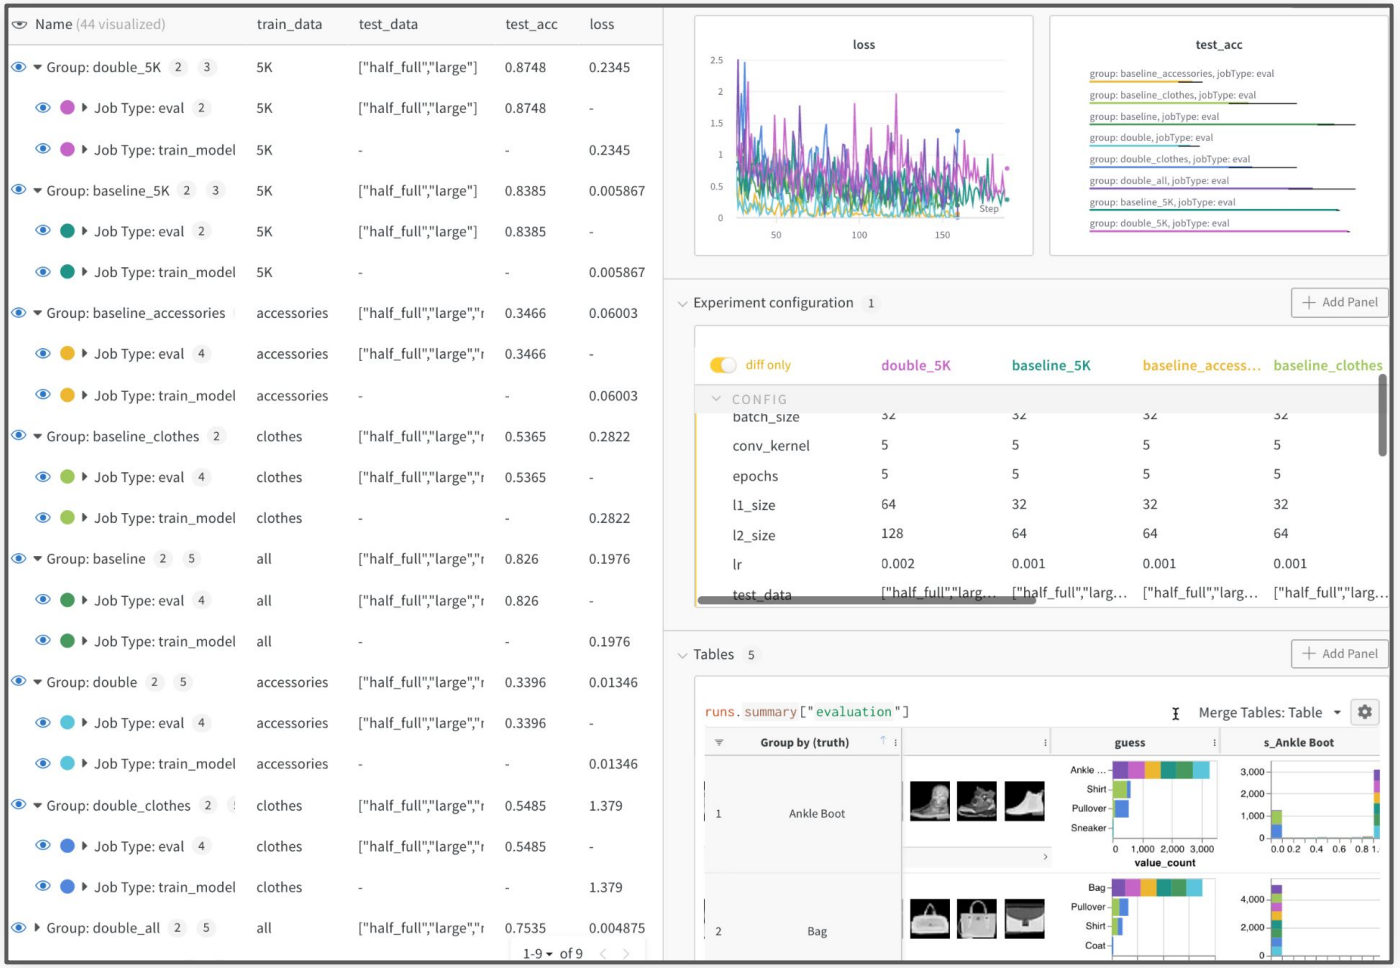
Name (44 visualized)
train_data
test_data
test_acc
loss
Group: double_5K 2 3 5K ["half_full","large"] 0.8748 0.2345
Job Type: eval 2 5K ["half_full","large"] 0.8748 -
Job Type: train_model 5K - - 0.2345
Group: baseline_5K 2 3 5K ["half_full","large"] 0.8385 0.005867
Job Type: eval 2 5K ["half_full","large"] 0.8385 -
Job Type: train_model 5K - - 0.005867
Group: baseline_accessories accessories ["half_full","large","r 0.3466 0.06003
Job Type: eval 4 accessories ["half_full","large","r 0.3466 -
Job Type: train_model accessories - - 0.06003
Group: baseline_clothes 2 clothes ["half_full","large","r 0.5365 0.2822
Job Type: eval 4 clothes ["half_full","large","r 0.5365 -
Job Type: train_model clothes - - 0.2822
Group: baseline 2 5 all ["half_full","large","r 0.826 0.1976
Job Type: eval 4 all ["half_full","large","r 0.826 -
Job Type: train_model all - - 0.1976
Group: double 2 5 accessories ["half_full","large","r 0.3396 0.01346
Job Type: eval 4 accessories ["half_full","large","r 0.3396 -
Job Type: train_model accessories - - 0.01346
Group: double_clothes 2 clothes ["half_full","large","r 0.5485 1.379
Job Type: eval 4 clothes ["half_full","large","r 0.5485 -
Job Type: train_model clothes - - 1.379
Group: double_all 2 5 all ["half_full","large","r 0.7535 0.004875
1-9 of 9
loss
test_acc
group: baseline_accessories, jobType: eval
group: baseline_clothes, jobType: eval
group: baseline, jobType: eval
group: double, jobType: eval
group: double_clothes, jobType: eval
group: double_all, jobType: eval
group: baseline_5K, jobType: eval
group: double_5K, jobType: eval
Experiment configuration 1
Add Panel
diff only
double_5K
baseline_5K
baseline_access...
baseline_clothes
CONFIG
batch_size
conv_kernel 5 5 5 5
epochs 5 5 5 5
l1_size 64 32 32 32
l2_size 128 64 64 64
lr 0.002 0.001 0.001 0.001
test_data
Tables 5
Add Panel
runs.summary["evaluation"]
Merge Tables: Table
Group by (truth)
guess
s_Ankle Boot
1 Ankle Boot
2 Bag
value_count

Check any metrics

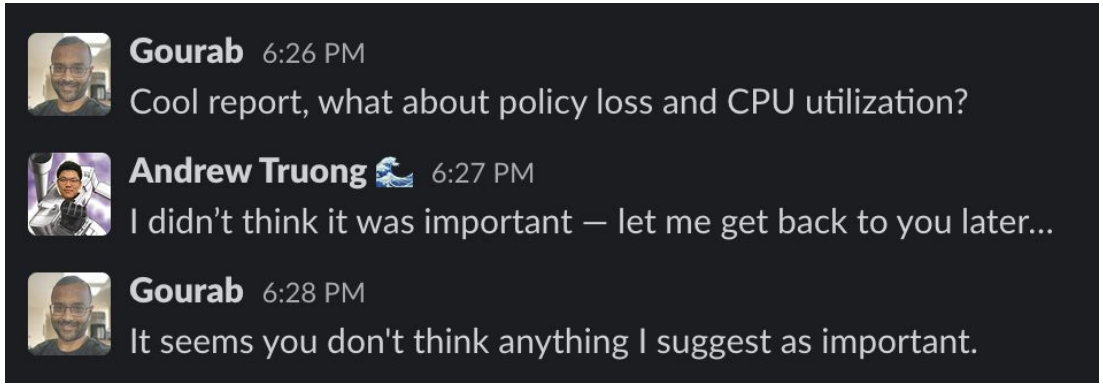

Interactive Dashboards

Check any metrics

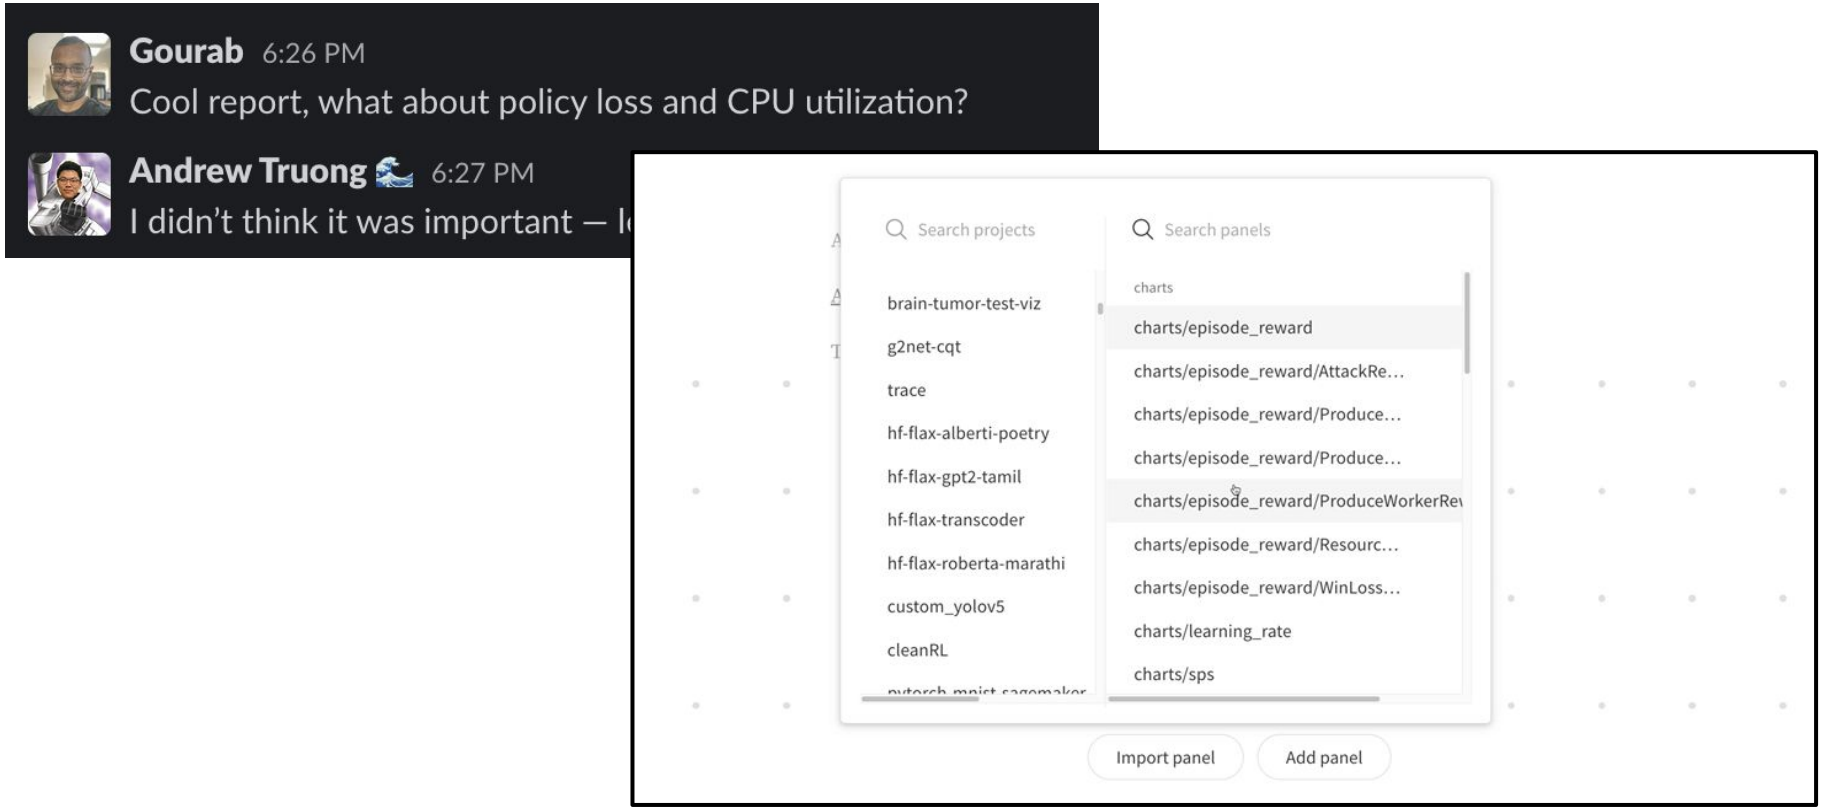

Unified Reporting and Dashboarding

Weights & Biases

Open Source Research

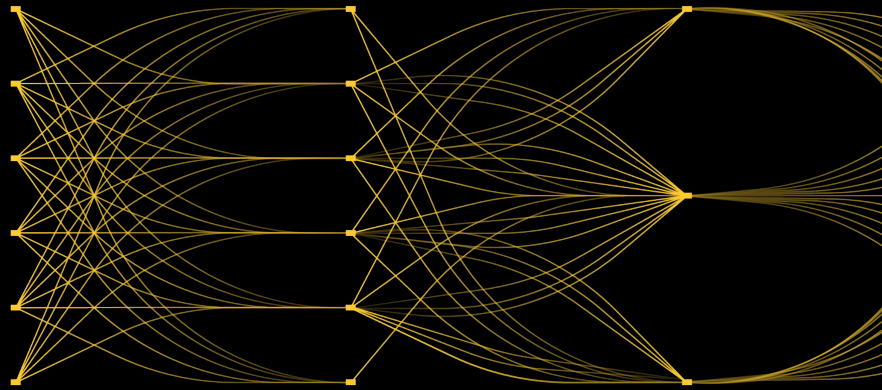

The world's leading ML teams trust us

COATUE INSIGHT BOND

figure eight

Trinity VENTURES

Felicis

The top open source research orgs use us

EleutherAI

LAION

CarperAI

Harmonai

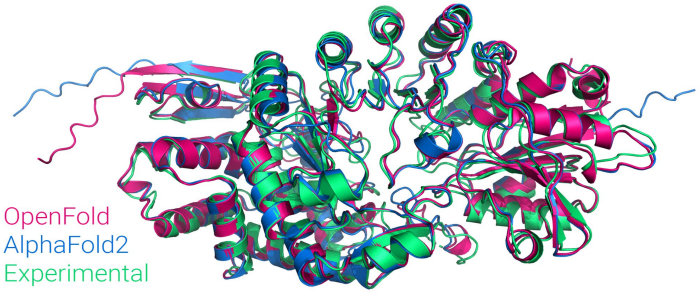

OpenFold

Craiyon

Weights & Biases

Integrations

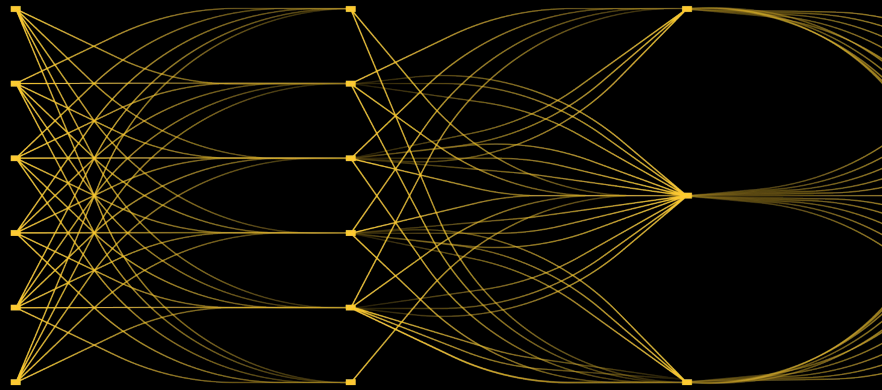

Fits into your **workflow**

Integrated into every popular ML framework

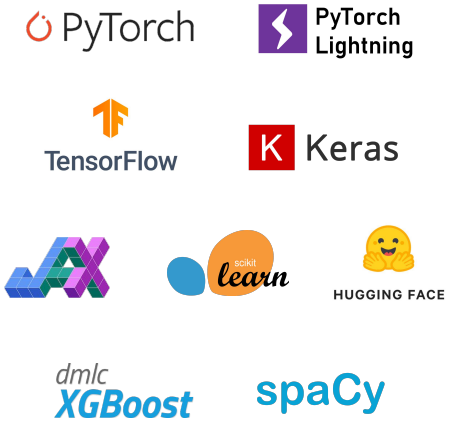

Instrumented into over 6,000 popular ML repos

Runs on every cloud or in your own infra

Weights & Biases

Competition Time!

wandb.me/xx

The W&B Course

www.wandb.courses

Weights & Biases

Thank You!

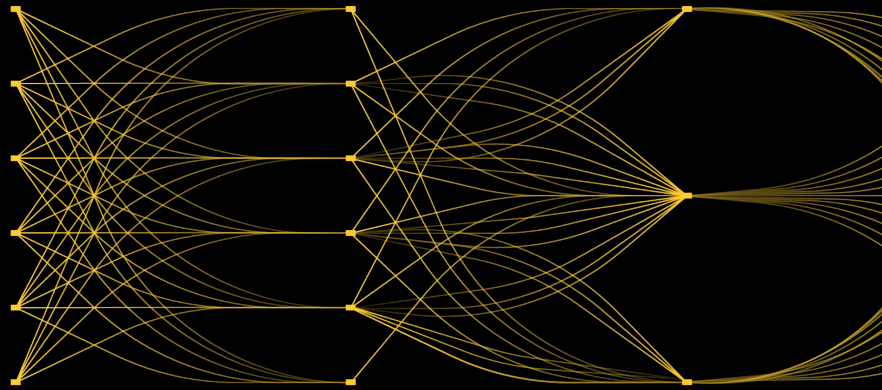

Join us and our ML community!

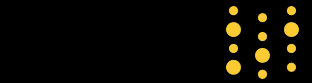

- Fully Connected – wandb.me/fc
- YouTube – wandb.me/youtube
- Twitter – wandb.me/twitter

Weights & Biases

Appendix

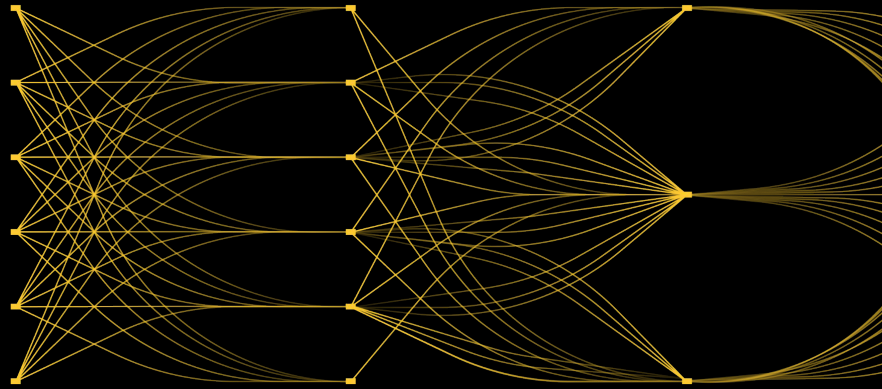

October 2022

Andrew’s Presentation

https://www.youtube.com/watch?v=Se1HvbAM0O4&t=12s