
Aim

Release 3.17.2

Gev Sogomonian, Gor Arakelyan et al.

Mar 28, 2023

CONTENTS

1	Overview	3
2	Get started with Aim	5
2.1	Installing Aim	5
2.2	Nightly releases	5
2.3	Initializing Aim repository	6
2.4	Tracking data with Aim SDK	6
2.5	Browsing results with Aim UI	7
2.6	Running Aim UI inside notebooks	7
2.7	Running Aim UI and tracking server inside Docker container	8
3	Integrate Aim into an existing project	9
3.1	Any python script	9
3.2	Integration with Pytorch Ignite	9
3.3	Integration with Pytorch Lightning	10
3.4	Integration with Hugging Face	11
3.5	Integration with Keras & tf.Keras	11
3.6	Integration with Keras Tuner	12
3.7	Integration with XGboost	12
3.8	Integration with CatBoost	12
3.9	Integration with LightGBM	13
3.10	Integration with fastai	13
3.11	Integration with MXNet	14
3.12	Integration with Optuna	14
3.13	Integration with PaddlePaddle	14
3.14	Integration with Stable-Baselines3	15
3.15	Integration with Acme	15
3.16	Integration with Prophet	16
3.17	What's next?	16
4	Migrate from other tools	17
4.1	Show TensorBoard logs in Aim	17
4.2	Show MLflow logs in Aim	19
4.3	Show Weights and Biases logs in Aim	19
5	Track media and objects	21
5.1	Tracking multiple values	21
5.2	Distribution tracking with Aim	22
5.3	Image tracking with Aim	22
5.4	Audio tracking with Aim	23

5.5	Text tracking with Aim	24
5.6	Figure tracking with Aim	25
5.7	Tracking matplotlib figures with Aim	25
5.8	Logging activeloop/hub dataset info with Aim	26
5.9	Log DVC metadata with Aim	27
5.10	Logging huggingface/datasets dataset info with Aim	28
6	Next steps	29
7	Overview	31
7.1	Pythonic Search	31
7.2	Explorers	31
7.3	Runs Management	32
7.4	Other Aim UI features	32
8	Home Page	33
8.1	Explore Aim	33
8.2	Overview	34
8.3	Latest Changes	35
9	Runs Management	37
9.1	Runs Explorer	37
9.2	Single run page	38
10	Explorers	43
10.1	Metrics Explorer	43
10.2	Images Explorer	49
10.3	Params Explorer	51
10.4	Scatters Explorer	52
11	Table	55
11.1	Overview	55
11.2	Delete the selected runs	55
11.3	Archive/Unarchive the selected runs	55
11.4	Hide/Show the selected runs	56
12	Bookmarks	57
12.1	Overview	57
12.2	The Bookmark Card	57
12.3	Delete Bookmark	58
13	Tags page	59
13.1	Overview	59
13.2	Create tag	59
13.3	Attach tag	60
13.4	Update attached tags	60
13.5	Edit tag	60
13.6	Delete tag	61
13.7	Used in overlay	61
14	Manage runs	63
14.1	Create runs	63
14.2	Continue runs	63
14.3	Delete runs	64
14.4	Upload runs	64

15	Configure runs	65
15.1	Reusing Run in your Repo	65
15.2	Defining custom location for your Repo	66
15.3	Organizing Runs in Experiments	66
15.4	Adding tags and params to Run	66
15.5	Training Run Reproducibility	67
15.6	Capturing Terminal logs	68
16	Query runs and objects	69
17	Query language basics	71
17.1	Introduction	71
17.2	Searching	71
17.3	Searching metrics and images	74
17.4	Security restrictions	75
18	Track experiments with aim remote server	77
18.1	Overview	77
18.2	Prerequisites	77
18.3	Scaling	77
18.4	Server-side setup	77
18.5	Client-side setup	78
18.6	SSL support	81
18.7	Message size limitations	81
18.8	Conclusion	81
19	Notify on failed/stuck runs	83
19.1	Run progress reporting SDK	83
20	Log messages during training process	85
20.1	The logging API	85
20.2	Notifying on logged messages	85
21	Set up the notification service	87
21.1	Quick start	87
21.2	Configuring notifiers	87
22	Define custom callbacks	91
22.1	Callbacks	91
22.2	Example	92
23	Host Aim on Kubernetes (K8S)	95
23.1	Dockerfile	95
23.2	Volume	96
23.3	Deployment	97
23.4	Service	98
23.5	Conclusion	99
24	Run Aim UI on Jupyter Notebook	101
25	Run Aim UI on SageMaker Notebook instance	103
25.1	Using Terminal	104
25.2	Using Notebook Extension	104
26	Integration guides	105
26.1	Pytorch Ignite	105

26.2	Pytorch Lightning	107
26.3	Hugging Face	107
26.4	TF/keras	107
26.5	XGBoost	108
26.6	Catboost	109
26.7	LightGBM	109
27	Overview	111
27.1	Aim Components	111
28	Data storage - where Aim data is collected	113
28.1	Storage structure	113
28.2	What is the index container?	114
28.3	How data written to/read from the storage?	114
29	Storage indexing - how Aim data is indexed	115
29.1	Background	115
29.2	How things worked before?	115
29.3	Automatic indexing	115
29.4	Conclusion	115
30	Concepts	117
30.1	Aim Run	117
30.2	Aim Repo	117
30.3	Run Params	117
30.4	Run Sequence	118
30.5	Sequence Context	118
31	Running Aim with profiling	119
31.1	Why would you need to enable the profiling.	119
32	Storing Metric sequence data for fast access	121
32.1	Background	121
32.2	What has changed?	121
32.3	What to do if the Metric has been logged already?	121
33	Track and compare GANs with Aim	123
33.1	Overview	123
33.2	Experiment	123
33.3	Track images with Aim	123
33.4	Explore the results with Aim UI	124
33.5	Conclusion	125
34	Aim CLI	127
34.1	init	127
34.2	version	128
34.3	up	128
34.4	server	128
34.5	runs	129
34.6	convert	130
34.7	storage	130
35	Aim status watcher CLI	133
35.1	start	133
35.2	notifiers	133

36 Aim SDK	135
36.1 aim.sdk.repo module	135
36.2 aim.sdk.run module	139
36.3 aim.sdk.objects.image	145
36.4 aim.sdk.objects.distribution	146
36.5 aim.sdk.objects.audio	147
36.6 aim.sdk.objects.text	148
36.7 aim.sdk.objects.figure	148
36.8 aim.sdk.sequence module	148
36.9 aim.sdk.sequences.metric module	149
36.10 aim.sdk.sequences.image_sequence module	149
36.11 aim.sdk.sequences.distribution_sequence module	149
36.12 aim.sdk.sequences.audio_sequence module	150
36.13 aim.sdk.sequences.text_sequence module	150
36.14 aim.sdk.sequences.figure_sequence module	150
36.15 aim.sdk.sequence_collection module	150
37 Aim Storage	153
37.1 aim.storage.arrayview module	153
38 Anonymized Telemetry	155
38.1 Motivation	155
38.2 How to opt out	155
39 Changelog	157
39.1 3.17.2 Mar 28, 2023	157
39.2 3.17.1 Mar 24, 2023	157
39.3 3.17.0 Mar 24, 2023	157
39.4 3.16.2 Mar 3, 2023	158
39.5 3.16.1 Feb 27, 2023	158
39.6 3.16.0 Feb 3, 2023	158
39.7 3.15.2 Dec 23, 2022	159
39.8 3.15.1 Dec 1, 2022	159
39.9 3.15.0 Nov 26, 2022	159
39.10 3.14.4 Nov 11, 2022	160
39.11 3.14.3 Oct 29, 2022	160
39.12 3.14.2 Oct 28, 2022	160
39.13 3.14.1 Oct 7, 2022	161
39.14 3.14.0 Oct 6, 2022	161
39.15 3.13.4 Sep 25, 2022	162
39.16 3.13.3 Sep 16, 2022	162
39.17 3.13.2 Sep 10, 2022	162
39.18 3.13.1 Sep 1, 2022	162
39.19 3.13.0 Aug 21, 2022	162
39.20 3.12.2 Aug 5, 2022	163
39.21 3.12.1 Aug 2, 2022	163
39.22 3.12.0 Jul 22, 2022	163
39.23 3.11.2 Jul 8, 2022	164
39.24 3.11.1 Jun 27, 2022	165
39.25 3.11.0 Jun 21, 2022	165
39.26 3.10.3 May 31, 2022	166
39.27 3.10.2 May 26, 2022	166
39.28 3.10.1 May 18, 2022	167
39.29 3.10.0 May 17, 2022	167

39.30 3.9.4 May 12, 2022	168
39.31 3.9.3 May 10, 2022	168
39.32 3.9.2 Apr 29, 2022	168
39.33 3.9.1 Apr 29, 2022	168
39.34 3.9.0 Apr 29, 2022	168
39.35 3.8.1 Apr 6, 2022	169
39.36 3.8.0 Mar 26, 2022	169
39.37 3.7.5 Mar 18, 2022	170
39.38 3.7.4 Mar 15, 2022	170
39.39 3.7.3 Mar 14, 2022	171
39.40 3.7.2 Mar 10, 2022	171
39.41 3.7.1 Mar 10, 2022	171
39.42 3.7.0 Mar 9, 2022	171
39.43 3.6.3 Mar 4, 2022	171
39.44 3.6.2 Mar 2, 2022	171
39.45 3.6.1 Feb 25, 2022	172
39.46 3.6.0 Feb 22 2022	172
39.47 3.5.4 Feb 15 2022	173
39.48 3.5.3 Feb 11 2022	173
39.49 3.5.2 Feb 10 2022	173
39.50 3.5.1 Feb 4 2022	173
39.51 3.5.0 Feb 3 2022	173
39.52 3.4.1 Jan 23 2022	174
39.53 3.4.0 Jan 22 2022	174
39.54 3.3.5 Jan 14 2022	175
39.55 3.3.4 Jan 10 2022	175
39.56 3.3.3 Dec 24 2021	175
39.57 3.3.2 Dec 20 2021	175
39.58 3.3.1 Dec 18 2021	175
39.59 3.3.0 Dec 17 2021	175
39.60 3.2.2 Dec 10 2021	176
39.61 3.2.1 Dec 8 2021	176
39.62 3.2.0 Dec 3 2021	176
39.63 3.1.1 Nov 25 2021	176
39.64 3.1.0 Nov 20 2021	176
39.65 3.0.7 Nov 17 2021	177
39.66 3.0.6 Nov 9 2021	177
39.67 3.0.5 Nov 9 2021	177
39.68 3.0.4 Nov 8 2021	177
39.69 3.0.3 Nov 4 2021	177
39.70 3.0.2 Oct 27 2021	177
39.71 3.0.1 Oct 22 2021	177
39.72 3.0.0 Oct 21 2021	178
39.73 2.7.1 Jun 30 2021	178
39.74 2.7.0 Jun 23 2021	178
39.75 2.6.0 Jun 12 2021	178
39.76 2.5.0 May 27 2021	178
39.77 2.4.0 May 13 2021	179
39.78 2.3.0 Apr 10 2021	179
39.79 2.2.1 Mar 31 2021	179
39.80 2.2.0 Mar 24 2021	179
39.81 2.1.6 Feb 26 2021	179
39.82 2.1.5 Jan 7 2021	179
39.83 2.1.4 Dec 2 2020	180

39.84 2.1.3 Nov 24 2020	180
39.85 2.1.2 Nov 24 2020	180
39.86 2.1.1 Nov 22 2020	180
39.87 2.1.0 Nov 19 2020	180
39.88 2.0.27 Nov 13 2020	180
39.89 2.0.26 Nov 10 2020	180
39.90 2.0.25 Nov 9 2020	181
39.91 2.0.24 Nov 8 2020	181
39.92 2.0.23 Nov 5 2020	181
39.93 2.0.22 Nov 5 2020	181
39.94 2.0.21 Nov 1 2020	181
39.95 2.0.20 Oct 26 2020	181
39.96 2.0.19 Oct 25 2020	181
39.97 2.0.18 Oct 7 2020	182
39.98 2.0.17 Oct 5 2020	182
39.99 2.0.16 Oct 2 2020	182
39.1002.0.15 Sep 21 2020	182
39.1012.0.14 Sep 21 2020	182
39.1022.0.13 Sep 21 2020	182
39.1032.0.12 Sep 12 2020	182
39.1042.0.11 Sep 11 2020	183
39.1052.0.10 Sep 10 2020	183
39.1062.0.9 Sep 10 2020	183
39.1072.0.8 Aug 26 2020	183
39.1082.0.7 Aug 21 2020	183
39.1092.0.6 Aug 13 2020	183
39.1102.0.5 Jul 18 2020	184
39.1112.0.4 Jul 18 2020	184
39.1122.0.3 Jul 8 2020	184
39.1132.0.2 Jul 7 2020	184
39.1142.0.1 Jun 24 2020	184
39.1152.0.0 Jun 18 2020	184
39.1161.2.17 May 8 2020	185
39.1171.2.16 Apr 29 2020	185
39.1181.2.15 Apr 29 2020	185
39.1191.2.14 Apr 27 2020	185
39.1201.2.13 Apr 25 2020	185
39.1211.2.12 Apr 20 2020	185
39.1221.2.11 Apr 16 2020	185
39.1231.2.10 Apr 16 2020	185
39.1241.2.9 Mar 16 2020	186
39.1251.2.8 Mar 15 2020	186
39.1261.2.7 Mar 14 2020	186
39.1271.2.6 Feb 28 2020	186
39.1281.2.5 Feb 25 2020	186
39.1291.2.4 Feb 20 2020	186
39.1301.2.3 Feb 13 2020	187
39.1311.2.2 Feb 13 2020	187
39.1321.2.1 Feb 13 2020	187
39.1331.2.0 Feb 13 2020	187
39.1341.1.1 Jan 14 2020	187
39.1351.1.0 Jan 12 2020	187
39.1361.0.2 Jan 7 2020	188
39.1371.0.1 Dec 26 2019	188

39.1381.0.0 Dec 25 2019	188
39.1390.2.9 Nov 30 2019	188
39.1400.2.8 Nov 30 2019	189
39.1410.2.7 Nov 14 2019	189
39.1420.2.6 Nov 5 2019	189
39.1430.2.5 Nov 4 2019	189
39.1440.2.4 Nov 3 2019	189
39.1450.2.3 Nov 3 2019	189
39.1460.2.2 Nov 1 2019	189
39.1470.2.1 Nov 1 2019	189
39.1480.2.0 Nov 1 2019	190
39.1490.1.0 - Sep 23 2019	190
40 Indices and tables	193
Python Module Index	195
Index	197

Jump to:

- [Quick start](#)
- [Aim UI](#)
- [Using Aim](#)
- [Understanding Aim](#)
- [Examples](#)
- [API reference](#)

If you have questions:

- [Open a feature request or report a bug](#)
- [Join our slack](#)

OVERVIEW

Jump to:

- [Quick start](#)
- [Aim UI](#)
- [Using Aim](#)
- [Understanding Aim](#)
- [Examples](#)
- [API reference](#)

If you have questions:

- [Open a feature request or report a bug](#)
- [Join our slack](#)

GET STARTED WITH AIM

This section shows a simple end-to-end aim setup. It starts from the installation, shows how to run Aim UI and explore the results. Use this as a starting point to get familiar with the basics of Aim while getting up and running.

2.1 Installing Aim

Aim is a python package available for Linux and MacOS for Python versions 3.6+. Install Aim using pip3:

```
pip3 install aim
```

Verify aim was properly installed

```
aim version
```

You should see the line listing newly installed version of Aim. For instance:

```
Aim v3.5.1
```

The installed package includes Python SDK needed for tracking training runs, UI for browsing the results and CLI for managing UI and results.

2.2 Nightly releases

Aim also provides daily dev packages with the features developed between main minor releases.

```
pip3 install --pre aim
```

Please note, that if the dependencies of aim are not already installed, this command will try to install the development versions of those packages as well.

Previous daily dev packages can be installed using the following command:

```
pip3 install aim==3.x.0.devyyyymmdd
```

[Release history](#)

2.3 Initializing Aim repository

Aim repository is the space where all your training runs are logged.

To initialize aim repo in the current working directory, run:

```
aim init
```

You should see something like this on your Command line:

```
Initialized a new Aim repository at /home/user/aim
```

Your workspace is now ready for tracking training runs with Aim.

2.4 Tracking data with Aim SDK

To start tracking, first create `aim.Run` object:

```
from aim import Run
run = Run()
```

`Run` class provides a dictionary-like interface for storing training hyperparameters and other dictionary-like metadata:

```
hparams_dict = {
    'learning_rate': 0.001,
    'batch_size': 32,
}
run['hparams'] = hparams_dict
```

These params can be used later on the UI to query runs, metrics, images. To track metrics with aim use the `Run.track` method:

```
run.track(3.0, name='loss')
```

The complete list of supported inputs is available in section “[Track media and objects](#)”

Here’s a full example demonstrating the steps above:

```
# aim_test.py
from aim import Run

run = Run()

# set training hyperparameters
run['hparams'] = {
    'learning_rate': 0.001,
    'batch_size': 32,
}

# log metric
for i in range(10):
    run.track(i, name='numbers')
```

Run the script above


```
python3 aim_test.py
```

Congrats! Your first run with Aim is ready! Now it is time to explore results with Aim UI.

2.5 Browsing results with Aim UI

Once the script above finishes you can open Aim UI and see the results:

```
aim up
```

You should see the following output meaning Aim UI is up and running:

```
Running Aim UI on repo `<Repo#-5930451821203570655 path=/.aim read_only=None>`
Open http://127.0.0.1:43800
Press Ctrl+C to exit
```

Open your browser and navigate to `http://127.0.0.1:43800` You should be able to see the home page of Aim UI:

Click on Metrics Explorer icon

In the Search bar select a “numbers” metric and click “Search”. You should be able to see line chart with tracked metric:

2.6 Running Aim UI inside notebooks

Run the following commands in the notebook to run the Aim UI:

1. Load Aim extension for notebooks:

```
%load_ext aim
```

1. Run `%aim up` to open Aim UI in the notebook:

```
%aim up
```

See [integration guide with Jupyter Notebook](#) for more details.

2.7 Running Aim UI and tracking server inside Docker container

Aim Docker [images](#) are available for running Aim UI and Aim Remote tracking server. Default command for Aim Docker image is `aim up`. To spin Docker container with Aim UI with port mapping:

```
docker run --publish 43800:43800 aimstack/aim
```

To run container with Aim Remote tracking server:

```
docker run --publish 53800:53800 aimstack/aim server
```

INTEGRATE AIM INTO AN EXISTING PROJECT

3.1 Any python script

```
from aim import Run

run = Run()

# Save inputs, hparams or any other `key: value` pairs
run['hparams'] = {
    'learning_rate': 0.001,
    'batch_size': 32,
}

# ...
for step in range(10):
    # Log metrics to visualize performance
    run.track(step, name='metric_name')
# ...
```

Aim easily integrates with your favourite ML frameworks.

Aim loggers give access to the `aim.Run` object instance via the `experiment` property. The `aim.Run` instance will help you to easily track additional metrics or set any other `key: value` pairs (params) relevant to your project.

In this way you can easily extend the default integrations. More info about this is available on [Integration guides section](#).

3.2 Integration with Pytorch Ignite

It only takes 2 steps to simply and easily inject Aim into pytorch ignite:

```
# import aim sdk designed for pytorch ignite
from aim.pytorch_ignite import AimLogger
```

Pytorch Ignite provides trainer objects to simplify the training process of pytorch model. We can attach the trainer object as `AimLogger`'s output handler to use the logger function defined by aim to simplify the process of tracking experiments. This process is divided into 2 steps:

Step 1. Create `AimLogger` object

```
aim_logger = AimLogger(  
    experiment='aim_on_pt_ignite',  
    train_metric_prefix='train_',  
    val_metric_prefix='val_',  
    test_metric_prefix='test_',  
)
```

Step 2. Attach output handler to the aim_logger object

```
aim_logger.attach_output_handler(  
    trainer,  
    event_name=Events.ITERATION_COMPLETED,  
    tag="train",  
    output_transform=lambda loss: {'loss': loss}  
)
```

Adapter source can be found [here](#). Example using Pytorch Ignite can be found [here](#).

3.3 Integration with Pytorch Lightning

We only require 2 steps to simply and easily inject Aim into pytorch lightning:

```
# import aim sdk designed for pl  
from aim.pytorch_lightning import AimLogger
```

Pytorch lighting provides trainer objects to simplify the training process of pytorch model. One of the parameters is called logger. We can use the logger function defined by aim to simplify the process of tracking experiments. This process is divided into 2 steps:

Step 1. Create AimLogger object

```
# track experimental data by using Aim  
aim_logger = AimLogger(  
    experiment='aim_on_pt_lightning',  
    train_metric_prefix='train_',  
    val_metric_prefix='val_',  
)
```

Step 2. Pass the aim_logger object as the logger argument

```
# track experimental data by using Aim  
trainer = Trainer(gpus=1, progress_bar_refresh_rate=20, max_epochs=5, logger=aim_logger)
```

Adapter source can be found [here](#). Example using Pytorch Lightning can be found [here](#).

3.4 Integration with Hugging Face

You only need 2 simple steps to employ Aim to collect data

Step 1: Import the sdk designed by Aim for Huggingface.

```
from aim.hugging_face import AimCallback
```

Step 2: Hugging Face has a trainer api to help us simplify the training process. This api provides a callback function to return the information that the user needs. Therefore, aim has specially designed SDK to simplify the process of the user writing callback functions, we only need to initialize AimCallback object as follows:

```
# Initialize aim_callback
aim_callback = AimCallback(experiment='huggingface_experiment')
# Initialize trainer
trainer = Trainer(
    model=model,
    args=training_args,
    train_dataset=small_train_dataset,
    eval_dataset=small_eval_dataset,
    compute_metrics=compute_metrics,
    callbacks=[aim_callback]
)
```

Adapter source can be found [here](#). Example using Hugging Face can be found [here](#).

3.5 Integration with Keras & tf.Keras

It only takes 2 steps to easily integrate aim in keras to record experimental information.

```
# call keras as the high api of tensorflow
from aim.tensorflow import AimCallback
# call keras library directly
from aim.keras import AimCallback
```

In keras, we call the `fit()` method of the model object to train the data. The callbacks are provided here. AimCallback inherits the usage specification of callbacks. We just need to add it to the callbacks list.

```
model.fit(x_train, y_train, epochs=5, callbacks=[
    # in case of tf.keras, we use aim.tensorflow.AimCallback
    AimCallback(experiment='aim_on_keras')
])
```

Adapter source can be found [here](#). Example using Keras can be found [here](#). Example using tf.Keras can be found [here](#).

3.6 Integration with Keras Tuner

It only takes 2 steps to easily integrate aim in keras to record experimental information.

```
from aim.keras_tuner import AimCallback
```

In kerastuner, we call the `search()` method of the tuner object to perform a search for best hyperparameter configurations. The callbacks are provided here. `AimCallback` inherits the usage specification of callbacks. We just need to add it to the callbacks list.

```
tuner.search(
    train_ds,
    validation_data=test_ds,
    callbacks=[AimCallback(tuner=tuner, repo='./aim_logs', experiment='keras_tuner_test
→')],
)
```

Adapter source can be found [here](#). Example using Keras Tuner can be found [here](#).

3.7 Integration with XGboost

Enjoy using aim to track xgboost experimental data which requires two simple steps:

Step 1: Explicitly import the `AimCallback` for tracking training data.

```
# call sdk aim.xgboost
from aim.xgboost import AimCallback
```

Step 2: XGboost provides the `xgboost.train` method for model training, in which the `callbacks` parameter can call back data information from the outside. Here we pass in `aimcallback` designed for tracking data information

```
xgboost.train(param, dtrain, num_round, watchlist,
              callbacks=[AimCallback(experiment='xgboost_test')])
```

Adapter source can be found [here](#). Example using XGboost can be found [here](#).

3.8 Integration with CatBoost

It only takes two steps to integrate Aim with `CatBoost`.

Step 1: Import `AimLogger` to track the training metadata.

```
# call sdk aim.catboost
from aim.catboost import AimLogger
```

Step 2: Pass the logger to the trainer.

Trainings in CatBoost are initiated with `fit` method. The method can be supplied with `log_cout` parameter to redirect output logs into a custom handler. Pass `AimLogger` to automatically track metrics and hyper-parameters with Aim. Depending on the training log output, an additional argument `logging_level` could be passed to make Catboost yield more logs to track test & best values.

```
model.fit(train_data, train_labels, log_cout=AimLogger(loss_function='Logloss'), logging_
↪ level='Info')
```

AimLogger also accepts `log_cout` parameter to preserve the default functionality of Catboost's log handling. You can pass your own handler, else it defaults to `sys.stdout`.

See AimLogger source [here](#). Check out a simple example with Aim and CatBoost [here](#).

3.9 Integration with LightGBM

Aim comes with a builtin callback designed to automatically track [LightGBM](#) trainings. It takes two steps to integrate Aim into your training script.

Step 1: Explicitly import the AimCallback for tracking training metadata.

```
from aim.lightgbm import AimCallback
```

Step 2: Pass the callback to callbacks list upon initiating your training.

```
gbm = lgb.train(params,
                lgb_train,
                num_boost_round=20,
                valid_sets=lgb_eval,
                callbacks=[AimCallback(experiment='lgb_test')])
```

While your training is running you can start `aim` up in another terminal session and observe the information in real time.

See AimCallback source [here](#). Check out a simple regression task example [here](#).

3.10 Integration with fastai

Aim comes with a builtin callback designed to automatically track [fastai](#) trainings. It takes two steps to integrate Aim into your training script.

Step 1: Explicitly import the AimCallback for tracking training metadata.

```
from aim.fastai import AimCallback
```

Step 2: Pass the callback to cbs list upon initiating your training.

```
learn = cnn_learner(dls, resnet18, pretrained=True,
                    loss_func=CrossEntropyLossFlat(),
                    metrics=accuracy, model_dir="/tmp/model/",
                    cbs=AimCallback(repo='.', experiment='fastai_example'))
```

See AimCallback source [here](#). Check out a simple regression task example [here](#).

3.11 Integration with MXNet

To track MXNet experiments use Aim callback designed for [MXNet](#) fit method. It takes two steps to integrate Aim into your training script.

Step 1: Import the AimLoggingHandler for tracking training metadata.

```
from aim.mxnet import AimLoggingHandler
```

Step 2: Pass a callback instance to event_handlers list upon initiating your training.

```
aim_log_handler = AimLoggingHandler(repo='.', experiment_name='mxnet_example',
                                   log_interval=1, metrics=[train_acc, train_loss, val_
↳ acc])

est.fit(train_data=train_data_loader, val_data=val_data_loader,
        epochs=num_epochs, event_handlers=[aim_log_handler])
```

See AimCallback source [here](#). Check out a simple regression task example [here](#).

3.12 Integration with Optuna

Aim provides a callback designed to automatically track [optuna](#) trainings. The `as_multirun` is a boolean argument. If `as_multirun` is set True then the callback will create a run for each trial. Otherwise it will track all of the results in a single run. One can also use the decorator function `track_in_aim` to log inside the objective function.

Step 1: Explicitly import the AimCallback for tracking training metadata.

```
from aim.optuna import AimCallback
```

Step 2: Pass the callback to cbs list upon initiating your training.

```
aim_callback = AimCallback(experiment_name="optuna_single_run")
study.optimize(objective, n_trials=10, callbacks=[aim_callback])
```

See AimCallback source [here](#). Check out a simple objective optimization example [here](#).

3.13 Integration with PaddlePaddle

Aim provides a built in callback to easily track [PaddlePaddle](#) trainings. It takes two steps to integrate Aim into your training script.

Step 1: Explicitly import the AimCallback for tracking training metadata.

```
from aim.paddle import AimCallback
```

Step 2: Pass the callback to callbacks list upon initiating your training.

```
callback = AimCallback(repo='.', experiment='paddle_test')
model.fit(train_dataset, eval_dataset, batch_size=64, callbacks=callback)
```

See AimCallback source [here](#). Check out a simple objective optimization example [here](#).

3.14 Integration with Stable-Baselines3

Aim provides a callback to easily track one of the reliable Reinforcement Learning implementations [Stable-Baselines3](#) trainings. It takes two steps to integrate Aim into your training script.

Step 1: Explicitly import the AimCallback for tracking training metadata.

```
from aim.sb3 import AimCallback
```

Step 2: Pass the callback to callback upon initiating your training.

```
model.learn(total_timesteps=10_000, callback=AimCallback(repo='.', experiment_name='sb3_
↳ test'))
```

See AimCallback source [here](#). Check out a simple objective optimization example [here](#).

3.15 Integration with Acme

Aim provides a built in callback to easily track [Acme](#) trainings. It takes few simple steps to integrate Aim into your training script.

Step 1: Explicitly import the AimCallback and AimWriter for tracking training metadata.

```
from aim.sdk.acme import AimCallback, AimWriter
```

Step 2: Initialize an Aim Run via AimCallback, and create a log factory using the Run.

```
aim_run = AimCallback(repo=".", experiment_name="acme_test")
def logger_factory(
    name: str,
    steps_key: Optional[str] = None,
    task_id: Optional[int] = None,
) -> loggers.Logger:
    return AimWriter(aim_run, name, steps_key, task_id)
```

Step 3: Pass the logger factory to logger_factory upon initiating your training.

```
experiment_config = experiments.ExperimentConfig(
    builder=d4pg_builder,
    environment_factory=make_environment,
    network_factory=network_factory,
    logger_factory=logger_factory,
    seed=0,
    max_num_actor_steps=5000)
```

See AimCallback source [here](#). Check out a simple objective optimization example [here](#).

3.16 Integration with Prophet

Aim provides an `AimLogger` object designed to track `Prophet` hyperparameters and metrics. It takes three steps to integrate Aim into your Prophet script.

Step 1: Explicitly import the `AimLogger`.

```
from aim.prophet import AimLogger
```

Step 2: After initializing a Prophet model, instantiate the `AimLogger` with your Prophet model.

```
model = Prophet()
logger = AimLogger(prophet_model=model, repo=".", experiment="prophet_test")
```

Step 3 (optional): pass any metrics you want after fitting the Prophet model.

```
metrics = {"backtest_mse": backtest_mse, "backtest_mape": backtest_mape}
logger.track_metrics(metrics)
```

Note that the metrics are assumed to be validation metrics by default. Alternatively, you can pass a `context` argument to the `track_metrics` method.

```
metrics = {"train_mse": backtest_mse, "train_mape": backtest_mape}
logger.track_metrics(metrics, context={"subset": "train"})
```

See `AimLogger` source [here](#). Check out a simple example [here](#).

3.17 What's next?

During the training process, you can start another terminal in the same directory, start `aim up` and you can observe the information in real time.

MIGRATE FROM OTHER TOOLS

The Aim explorers add true superpowers to the AI engineer's arsenal. However not all training runs may have been tracked by Aim. So it is important to be able to port existing training run logs. There might be 1000s of training runs tracked with other tools. Aim has built-in converters to easily migrate logs from other tools. These migrations cover the most common usage scenarios. In case of custom and complex scenarios you can use Aim SDK to implement your own conversion script.

As of Aim v3.6.0 the following converters are supported:

- *TensorBoard logs converter*
- *MLFlow logs converter*
- *Weights & Biases logs converter*

We are working to constantly improve existing converters and implement new ones.

4.1 Show TensorBoard logs in Aim

Aim gives you a possibility to convert **TensorBoard** logs into native format and show them directly inside the Aim UI.

Before showing the logs in Aim, the files have to pass the conversion process.

Please note that only the following plugins are currently supported

- scalar
- image

To convert TensorBoard logs, `aim convert` command must be run on your log directory.

```
aim convert tensorboard --logdir ~/my_logdir
```

To make conversion process smooth please ensure that logs directory structure follows conventions below. Consider the following directory hierarchy:

```
~/my_logdir/  
-> run_1/  
    |> <evnet_file_1>  
    |> <evnet_file_2>  
-> group_1/  
    |> <evnet_file_3> (THIS LOG WILL BE IGNORED)  
    |> run_2/  
        |> train/  
            |> <evnet_file_4>
```

(continues on next page)

(continued from previous page)

```

└─> <evnet_file_5>
    └─> validate/
        ├──> <evnet_file_6>
        ├──> <evnet_file_7>
        ├──> <evnet_file_8> (IGNORED IF "--flat" IS ACTIVE)
        └──> <evnet_file_9> (IGNORED IF "--flat" IS ACTIVE)
└─> run_3/
    ├──> <evnet_file_10>
    └──> <evnet_file_11>
-> <evnet_file_12> (THIS LOG WILL BE IGNORED)
-> <evnet_file_13> (THIS LOG WILL BE IGNORED)

```

Note that directory naming is not mandated and its up to you how to name them.

The conversion logic categorizes your hierarchy into one of group, run and context categories where.

- **group**: Is a directory which has one or more **run** directories inside it,
- **run**: Is a directory which has either log files or context directory inside it,
- **context**: Is a directory inside of **run** directory which has at least one log file inside it.

Conversion process will scan and determine `run` directories for your hierarchy and will create a distinct run for each of them.

From the hierarchy example above you can see that the following log files will be ignored since the converter treats them as unorganized log files.

- ~/my_logdir/group_1/evnet_file_3
- ~/my_logdir/evnet_file_12
- ~/my_logdir/evnet_file_13

All other logs will either have `Context` or `No Context`. Context of the log is the name of the parent directory if the parent directory hasn't been categorized into neither as `run` nor `group` category.

For example:

- Log files right underneath `run_1`, `run_2` and `run_3` will have no context
- Log files in `run_2/train` and `run_2/validate` will have `train` and `validate` as context accordingly.

In case the converter finds unorganized log files in your hierarchy a warning message will be issued.

To make the converter process these logs, consider re-structuring your directories so that it matches the sample structure. (i.e. create a new directory and moving your unorganized logs there)

You can make converter treat every directory as a distinct run by supplying `--flat` option. In this case the following directories will be categorized as a `run` directory.

- ~/my_logdir/run_1/
- ~/my_logdir/group_1/run_2/train/
- ~/my_logdir/group_1/run_2/validate/
- ~/my_logdir/group_1/run_3/

The log files in all other directories will be ignored.

4.2 Show MLflow logs in Aim

Aim gives you a possibility to convert [MLflow](#) runs into native format and show them directly on Aim UI.

Before showing your MLflow runs on Aim, they need to pass conversion process where your metrics, tags, parameters, run description/notes and *some* artifacts will be transferred into Aim storage.

Please note that as for now, only the artifacts having the following file extensions will be transferred into Aim storage!

- Images: ('jpg', 'bmp', 'jpeg', 'png', 'gif', 'svg')
- Texts: ('txt', 'log', 'py', 'js', 'yaml', 'yml', 'json', 'csv', 'tsv', 'md', 'rst', 'jsonnet')
- Sound/Audios: ('flac', 'mp3', 'wav')

To convert MLflow runs, `aim convert mlflow` command must be run on your log directory:

```
$ aim init
$ aim convert mlflow --tracking_uri 'file:///Users/aim_user/mlruns'
```

You can also set the `MLFLOW_TRACKING_URI` environment variable to have MLflow find a URI from there. In both cases, the URI can either be an HTTP/HTTPS URI for a remote server, a database connection string, or a local path to log data to a directory.

The conversion process will iterate over all your Experiments and create a distinct run for each run inside the experiment. If you want to process only a single experiment, you can provide the experiment id or name to the conversion command:

```
$ aim convert mlflow --tracking_uri 'file:///Users/aim_user/mlruns' --experiment 0
```

While converting the artifacts, the converter will try to determine file content type only based on its extension. A warning message will be issued if artifact cannot be categorized, these artifacts will not be transferred to aim! Please check the command output logs if you fail to see your artifact in Aim's web.

4.3 Show Weights and Biases logs in Aim

Aim gives you a possibility to convert [Weights & Biases](#) runs into native format and explore them via Aim UI.

To be able to explore Weights & Biases runs with Aim, please run the WandB to Aim converter. All the metrics, tags, config, experiment description/notes and *some* artifacts will be converted and stored in `.aim` repo.

Execute `aim convert wandb` command to start converting Weights & Biases experiments:

```
$ aim init
$ aim convert wandb --entity 'my_team' --project 'test_project'
```

The convertor will iterate over all the experiments and create a distinct Aim run for each experiment. If you want to process only a single experiment, please specify the experiment id or name when running the command:

```
$ aim convert wandb --entity my_team --project test_project --run-id '<HASH_OF_RUN>'
```

While converting the artifacts, the converter will try to determine file content type based on its extension. A warning message will be issued if artifact will not be categorized, these artifacts will not be transferred to aim. Please note that current implementation tracks only the primitive datatypes.

If there are any issues with the conversion process, please [submit an issue](#).

Once conversion process is complete - you can enjoy the power of Aim

TRACK MEDIA AND OBJECTS

Aim supports variety of data sources. Basic logging of Run params covers Python builtin types (such as `int`, `float`, `bool`, `bytes` and `str`) as well as composition of those into dictionaries, lists, tuples at any depth.

In addition to the builtin types, Aim provides native support for [OmegaConf](#) configs, thus simplifying integration for projects running with [Hydra](#).

Starting from v3.6.0 Aim provides integration with [activeLoop/hub](#) datasets. Hub is the open-source dataset format for AI.

Tracking of data includes metrics, images, audio, text and chart figures. Here's the complete list of Aim objects provided by the package:

- Metrics
- *Distribution*
- *Image*
- *Audio*
- *Text*
- *Figure*

5.1 Tracking multiple values

Starting from v3.14.0 Aim provides ability to track multiple values at once. To track multiple values of different sequences of the same context, just pass a dictionary of `<name: value>` pairs to the track method:

```
aim_run.track({'accuracy': 98.2, 'loss': 0.001}, context={'subset': 'train'}, step=10,   
↳ epoch=1)
```

Note: The name argument should be set to `None` in this case.

More examples

The following code snippet

```
from aim import Run

aim_run = Run()

metrics = {'accuracy': 0.72, 'f1': 0.99}
for metric, val in metrics.items():
    aim_run.track(val, name=metric, context={'subset': 'train'})
```

can be simplified to:

```
from aim import Run

aim_run = Run()

aim_run.track({'accuracy': 0.72, 'f1': 0.99}, context={'subset': 'train'})
```

5.2 Distribution tracking with Aim

You can store distribution objects in Aim repository using our `aim.Distribution` object.

```
from aim import Distribution
```

`aim.Distribution` accepts the following parameters

- **distribution**: array-like object used to construct `aim.Distribution`.
- **bin_count**: Optional distribution bin count. 64 by default, max 512.

Your data is converted to `numpy.histogram` upon initialization of the object.

Simple example of initializing and tracking distribution

```
import random
from aim import Run, Distribution

run = Run()
d = Distribution(
    distribution=[random.randrange(0, 10000) for _ in range(1000)],
    bin_count=250
)
run.track(d, name='dist', step=0)
```

5.3 Image tracking with Aim

Aim lets you track an image using `aim.Image` object

To get started, first import the `Image` object into your code.

```
from aim import Image
```

Our `Image` object uses `Pillow` under the hood. `Image` object supports the following inputs as data source.

- Path to an image file
- **PIL** (Pillow object)
- **torch.Tensor** (PyTorch tensor object)
- **tf.Tensor** (TensorFlow tensor object)
- **np.array** (Numpy array object)
- **matplotlib.figure.Figure** (matplotlib figure object)

Here's an example of tracking image from file path

```
path = "~/test_image.png"
aim_image = Image(path)
```

Image object also has the following arguments:

```
caption (:obj: `str`, optional): Optional image caption. '' by default.
format (:obj: `str`, optional): Parameter for PIL's .save() method. 'png' by default.
quality (:obj: `int`, optional): Parameter for PIL's .save() method. 85 by default.
optimize (:obj: `bool`, optional): Parameter for PIL's .save() method. False by default.
```

For more information on the format, quality and optimize parameters, refer to [Pillow documentation](#).

Using these parameters you can manipulate image quality and/or convert the image format from .png to jpeg or to any other format (which is supported by Pillow)

```
from aim import Run, Image

# Initialize a new run
run = Run()

for step in range(1000):
    # Log image
    path = f"~/test_image_{step}.png"
    aim_image = Image(
        path,
        format='jpeg',
        optimize=True,
        quality=50
    )

    run.track(aim_image, name='images', step=step)
```

5.4 Audio tracking with Aim

Aim lets your track an audio data using `aim.Audio` object

To get started, first import the Audio object into your code.

```
from aim import Audio
```

You can use Audio object to track MP3, WAV and FLAC audio data. Audio object supports the following data as input.

- File path
- Raw bytes
- `io.BytesIO` stream
- Numpy array (only for WAV audio format)

This object comes with the following optional arguments.

```
format (:obj:`str`): Format of the audio source. Choices are ('flac', 'mp3', 'wav')
rate    (:obj:`int`): Only for WAV. Rate of the audio file, defaults to 22500
caption (:obj:`str`): Optional audio caption. An empty string by default.
```

Complete example of tracking WAV audio data.

```
import os.path
from aim import Run, Audio

# Initialize a new run
run = Run()

for step in range(1000):
    # Log image
    path = f"~/test_audio_{step}.mp3"
    aim_audio = Audio(
        path,
        format='mp3',
        caption=os.path.basename(path)
    )

    run.track(aim_audio, name='audios', step=step)
```

5.5 Text tracking with Aim

Aim lets you track text/string during your training process.

To get started, first import the Text object into your code.

```
from aim import Text
```

In order to use the Text object, you just need to ensure that your input data type is a string.

Here's an example of Text usage:

```
import random
import string
from aim import Run, Text

# Initialize a new run
run = Run()

for step in range(100):
    # Generate a random string for this example
    random_str = ''.join(random.choices(
        string.ascii_uppercase +
        string.digits, k=20)
    )
    aim_text = Text(random_str)
    run.track(aim_text, name='text', step=step)
```

5.6 Figure tracking with Aim

Aim provides a Figure object which can be used to track `plotly` and `matplotlib` figures.

To get started, first import the Figure object into your code.

```
from aim import Figure
```

You should pass either `Plotly Figure` or `matplotlib Figure` as input source to Aim's Figure object.

Here's an example of tracking a plotly figure

```
import plotly.express as px
from aim import Run, Figure

# Initialize a new run
run = Run()

# First we create Plotly figure
fig = px.bar(x=["a", "b", "c"], y=[1, 3, 2])

# Now we convert it to Aim Figure
aim_figure = Figure(fig)

run.track(aim_figure, name="plotly_figures", step=0)
```

It is also easy to track `matplotlib` figure. Please note that the conversion process is done by Plotly under the hood.

```
from aim import Run, Figure
import matplotlib.pyplot as plt

# Initialize a new run
run = Run()

# define matplotlib figure
fig = plt.figure()
plt.plot([1, 2, 3])
plt.close(fig)

# Now we convert it to Aim Figure using (Plotly's functions)
aim_figure = Figure(fig)

run.track(aim_figure, name="matplotlib_figures", step=0)
```

5.7 Tracking matplotlib figures with Aim

In order to track `matplotlib` figures with Aim, either pass the `matplotlib` figure to Aim's Image or Figure object.

5.7.1 Converting matplotlib to Aim Image

```
from aim import Run, Image
import matplotlib.pyplot as plt

run = Run()

# define matplotlib figure
fig = plt.figure()
plt.plot([1, 2, 3])
plt.close(fig)

# pass it to aim Image
aim_img = Image(fig)
run.track(aim_img, step=0, name="matplotlib_images")
```

5.7.2 Converting matplotlib to Aim Figure

Please note that the conversion process is done by [Plotly](#) under the hood.

```
from aim import Run, Figure
import matplotlib.pyplot as plt

run = Run()

# define matplotlib figure
fig = plt.figure()
plt.plot([1, 2, 3])
plt.close(fig)

aim_figure = Figure(fig)
run.track(aim_figure, step=0, name="matplotlib_figures")
```

5.8 Logging activeloop/hub dataset info with Aim

Aim provides wrapper object for `hub.dataset`. It allows to store the dataset info as a Run parameter and retrieve it later just as any other Run param. Here is an example of using Aim to log dataset info:

```
import hub

from aim.sdk.objects.plugins.hub_dataset import HubDataset
from aim.sdk import Run

# create dataset object
ds = hub.dataset('hub://activeloop/cifar100-test')

# log dataset metadata
run = Run(system_tracking_interval=None)
run['hub_ds'] = HubDataset(ds)
```

5.9 Log DVC metadata with Aim

If you are using [DVC](#) to version your datasets or track checkpoints / other large chunks of data, you can use Aim to record the info about the tracked files and datasets on Aim. This will allow to easily connect your datasets info to the tracked experiments. Here is how the code looks like

```
from aim.sdk import Run
from aim.sdk.objects.plugins.dvc_metadata import DvcData

run = Run(system_tracking_interval=None)

path_to_dvc_repo = '.'
run['dvc_info'] = DvcData(path_to_dvc_repo)
```

If we consider the following [sample repo](#) provided by DVC team:

Run the following command to list repository contents, including files and directories tracked by DVC and by Git.

```
$ git clone https://github.com/iterative/example-get-started
$ cd example-get-started
$ dvc list .
.dvcignore
.github
.gitignore
README.md
data
dvc.lock
dvc.yaml
model.pkl
params.yaml
prc.json
roc.json
scores.json
src
```

If we apply our previous code snippet on the same repo - we can observe the same information added to Run parameters.

```
{
  'dvc_info.dataset.source': 'dvc',
  'dvc_info.dataset.tracked_files': [
    '.dvcignore', '.github', '.gitignore',
    'README.md', 'data', 'dvc.lock',
    'dvc.yaml', 'model.pkl', 'params.yaml',
    'prc.json', 'roc.json', 'scores.json', 'src'
  ]
}
```

5.10 Logging huggingface/datasets dataset info with Aim

Aim provides wrapper object for datasets. It allows to store the dataset info as a Run parameter and retrieve it later just as any other Run param. Here is an example of using Aim to log dataset info:

```
from datasets import load_dataset

from aim import Run
from aim.hf_dataset import HFDataset

# create dataset object
dataset = load_dataset('rotten_tomatoes')

# store dataset metadata
run = Run()
run['datasets_info'] = HFDataset(dataset)
```

NEXT STEPS

In the Quick Start section you learned the basics of Aim functionality and how to use the Aim SDK.

In the following sections we will cover different aspects of Aim in more detail. Aim is designed and built to explore and make sense of large volumes of training run logs.

If you are interested in exploring more about Aim, please navigate to [Using Aim](#) section.

You can browse real-life Aim use-cases in [Examples](#) section.

If you want to learn more about how Aim works under the hood and get familiar with Aim terminology, navigate to [Understanding Aim](#) section.

Detailed API descriptions as well as CLI commands available at [References](#).

OVERVIEW

Aim UI is one of the key two interfaces to interact with the ML training runs tracked by Aim. It's super-powerful for comparing large number of experiments.

The Aim UI is mainly built around the Explorers that allow to query and compare metrics, images and params, scatterplots etc.

Explorers are built around the way to query runs and compare them. Each explorer has a way to navigate to a single run where all the run-related information is visualized and can be used to do deep-dive into runs.

Besides these, Aim UI also allows to tag the runs, delete/archive them and save Explorers state to share with the team.

7.1 Pythonic Search

Once you track your experiments with Aim, they are available to be searched. There are three objects that could be searched:

- Runs
- Metrics
- Images

This basically means that you can write a python if-statement over everything you have tracked with a few options of the output.

Explorers encompass these outputs and give you unique super-powers to manipulate and compare the results of the search.

7.2 Explorers

Explorers are powerful tools to query ML training runs, select specific type of data tracked (metrics, params etc), apply modifications to them and compare them.

The runs are compared by grouping and dividing them via the tracked hyperparams. **Every** param (even the system params) are available on Aim - to be used for runs comparison on all Explorers.

7.3 Runs Management

Runs Management includes the Runs Explorer. You can search through runs on the Runs Explorer and have the holistic view of your runs, their diff, the last values of the metrics, all the params etc.

This also contains the Single Run Page that will help you observe everything you have tracked for your run - all in one place. This includes params, metrics, images, distributions etc.

7.4 Other Aim UI features

Besides this, you can also save the Explorer states for reproducible experiment analysis. Aim also enables ways to tag the runs.

HOME PAGE

Aim Home Page is a high-level overview of your training activities and how to find your way around the Aim UI.

There are three main sections:

- *Explore Aim*
 - *Overview*
 - *Latest Changes*
-

8.1 Explore Aim

In the Explore Aim section, you will find a list of recommended resources to help you get started with Aim or navigate through an existing Aim project. Here are its main sections

- *Quick navigation to some of the most used features queries*
- *Recent searches*
- *Experiments table*
- *Tags table*
- *Bookmarks*

8.1.1 Quick Navigation

The quick navigation section provides a quick way to navigate through your runs. It includes:

- Active Runs
- Archived Runs
- Last week's runs

8.1.2 Recent searches

In this block, you will find your last three successful queries with the ability to navigate through them.

8.1.3 Experiments table

This block displays your experiment data in the table format. Also, it provides the ability to compare the selected ones on the explorer pages.

8.1.4 Tags table

The tags table provides a quick way to navigate to your tags and compare them on the other explorer pages.

8.1.5 Bookmarks

This section displays your bookmarks and provides the ability to navigate through them.

8.2 Overview

The overview section is a high-level overview of your training activities. It includes:

- *Statistics of your training activities*
- *Active Runs table*
- *Contributions heatmap*
- *Activity feed*

8.2.1 Statistics

Here you can find the total number of runs and tracked sequences. Also, each card is clickable and you can navigate to the Runs and Sequences pages.

8.2.2 Active Runs table

The active runs table provides the list of your active runs. Also, the ability to compare your runs in other explorer pages.

8.2.3 Contributions heatmap

The heatmap shows the intensity of experiments you have made for the day - the darker the color the more experiments. Each cell represents the set of training runs for that day. The cells are clickable! Once clicked, you will navigate to the Runs explorer page and automatically query the runs made on that day.

8.2.4 Activity feed

This section represents the activity feed of your contributions. It includes the following information:

- The date of the activity
 - The type of activity
 - The name of the activity
-

8.3 Latest Changes

In this section, you can find the latest changes in Aim. It includes:

- The latest changes to the Aim version
- The Changelog of the Aim versions
- The helpful doc links for exploring the Aim

RUNS MANAGEMENT

9.1 Runs Explorer

9.1.1 Overview

To navigate into Runs Explorer, click on the Runs navigation item from the left sidebar.

Runs explorer helps you to

- *Search runs with pythonic query*
- *Observe runs in real time*
- *Delete or archive runs*
- *Export Runs report*

Search Runs

Use Search bar to query runs with [Aim QL](#).

Follow runs in real time

Switch **Live Update** to turn on the real time mode.

Delete or archive runs

Step 1: Select runs on the runs table:

Click on the **Archive** button. Confirmation popup appears. Click **Archive** again and the runs are archived!

In order to batch delete the selected runs, just use the **Delete** button. In this case as well, press **Delete** again on the confirmation popup and the runs will be hard deleted. ***Warning:** this operation is irreversible and the runs are deleted from the disk.

Export Runs report

Generate Runs CSV report by clicking on the **Export** button on the Runs table.

9.2 Single run page

Each training run has a dedicated page on Aim. Use the single run page to observe all the tracked metadata associated with that run.

Here are the tabs available on the single run page. Each tab visualizes respective tracked metadata or empty if not tracked.

- *Overview*
- *Params*
- *Metrics*
- *System*
- *Distributions*
- *Images*
- *Audios*
- *Texts*
- *Figures*
- *Terminal Logs*
- *Settings*

9.2.1 Overview

Overview tab shows overall info about the run.

- These cards can contain information about Parameters, Metrics, System Metrics, CLI Arguments, Environment Variables, Packages and Git information. With this data, you can easily reproduce your run.
- Sidebar contains information about Run Date, Run Duration, Run Hash, attached Tags and gives the ability to navigate through tabs.

Also, you can apply advanced searching/filtering to those card tables.

9.2.2 Params

Params tab contains a JSON-like visualization of all the tracked params data related to a single run of interest.

9.2.3 Metrics

Metrics tab contains the visualizations of all the metrics tracked for the given run.

When exploring run results, some metrics are looked at more often. Pin those metrics to the top of the page to find them easier and faster. Undo pinning by clicking on “unpin” button.

Note: you can track arbitrary number of runs with lots of steps with Aim!

9.2.4 System

Aim automatically tracks system metrics, so you can use them in order to detect potential resource mismanagements or anomalies.

System tab contains all the tracked system metrics for a single run.

9.2.5 Distributions

You can track the gradient, the weights and the biases distributions of all the layers for lots of steps with Aim.

The Distributions tab will allow you to observe them for a single run. You can also

- navigate between the layers
- search for distribution on specific steps

The single run distributions tab is quite powerful!

9.2.6 Images

The Images tab contains all the tracked images of a single run. You can track runs with different contexts and at different steps of training.

On the left-hand side are the names of different image-sets you have tracked along with their `context` unpacked.

Usually the images are tracked at diff steps and with batches. This control will allow you to quickly slice and dice the specific subset of images to view. Use these sliders to search

- which subset of steps you'd like to see (on the left-hand side)
- which indices you'd like to see (on the right-hand side)

If there is only 1 step or only 1 index, you will see the info message instead of the control

9.2.7 Audios

You can track audios with Aim. Use the Audios tab to view and play the tracked audios of a single run.

9.2.8 Texts

Use the Texts tab to view and search all the texts tracked for a single run.

On the left-hand side you will see the name and context of the tracked texts. You can use the search-bar on top to search the text with regexp or just match word or case.

Use the bottom controllers to control the steps and the indices of the tracked texts too.

9.2.9 Figures

Aim allows tracking Plotly and matplotlib figures. On the Figures tab you can view all the track figures over different contexts and steps.

Note: Aim will render figures with passed or default dimensions. There will be scrolls if the size exceeds the plotly container space of the Figures tab.

9.2.10 Terminal Logs

Aim streams the process output logs (terminal logs) to the Aim UI automatically near-real-time. The terminal logs are displayed under Logs tab on the single run page. The terminal logs can be disabled programmatically if needed. [More on terminal logs here](#).

9.2.11 Settings

Use the Settings tab to delete or archive the single run

Delete Run

Archive Run

EXPLORERS

Explorers will help you to compare 1000s of AI experiments with a few clicks. Explorers are the main tools that Aim is built around.

In this section we will go through the Aim explorers, introduce their features and how to use them.

10.1 Metrics Explorer

10.1.1 Overview

Use Metrics explorer to search and compare 1000s of ML training metrics.

The Metrics Explorer allows you to search, group and compare your metrics. Due to this and number of other visual features on the Metrics Explorer, you will save considerable amounts of time when comparing experiments- compared to other open-source experiment tracking tools.

The Metric Explorer has the following main sections:

- **Metrics Select:** to select the metrics for exploration
- **Search bar:** to query the runs for exploration
- **Charts explorer:** the space where the metrics are rendered
- **Metrics modifiers:** all the grouping, chart division and other metrics modifier tools
- **Context table:** the full information about the selected metrics is available

There is also an advanced mode of search is available too where you can use the full Aim QL (more on this further in this section).

There are two ways you can query metrics and runs

- *Select metrics and query runs*
- *Advanced Search mode*

An overview of what you can do with queried metrics - the modifiers:

- *Group by tracked parameters*
- *Aggregate grouped metrics*
- *Configure axes properties (x-axis alignment, axes range)*
- *Change scale of the axes (linear or log)*
- *Apply smoothing*
- *Ignore outliers*

- *Metric highlight modes (metric on hover, run on hover)*
- *Set chart tooltip parameters*
- *Apply zoom in/out on charts*
- *Export chart as image*

10.1.2 Select metrics and query runs

On the Metrics Explorer, there is + **Metrics** button. Once pressed, a dropdown will appear with all your tracked metrics and their contexts flattened. The dropdown is searchable - so you can get to your metric of interest within a stroke!

The Search bar is located below the + **Metric** button. It allows to do a pythonic query (that is eval-ed as python statement) over every param you have tracked.

Search runs with [Aim QL](#)

10.1.3 Advanced Search mode

Once you press the **Enable advanced search mode** button underneath the main **Search** button, it will enable the full Aim QL search editor - to query the metrics, the runs via full Aim QL

Here is an example:

```
((metric.name == 'bleu' and metric.context.subset == 'val') or (metric.name == 'loss' and metric.context.subset == 'train')) and 1e-5 < run.hparams.learning_rate < 1e-2
```

10.1.4 Group by any parameter

Grouping selected metrics by any tracked params will allow you to quickly distinguish the most impactful params, decisions you have made (the preprocessing steps, the hyperparams etc).

The parameters include not only the ones you have tracked but also the native Aim objects too such as

- `metric.name`
- `metric.context.[context_key]`
- `run.hash`

There are several ways you can group the selected metrics and runs - by color, by stroke and by chart.

Group by Color

Use this to divide the selected metrics into different clusters based on selected values of params. Each cluster gets colored differently.

There are a number of options available when grouping

- **group by values** - divides into clusters as per the values of selected params)
- **reverse grouping** - divides into clusters by every param except for the chosen one.

The grouping colors are picked randomly, however it is possible to fix with the advanced coloring features.

Here are the features in the advanced mode:

- Fix the colors of the grouping
- Control the color palette to use during the grouping

Group by Stroke

Groups the metrics by a stroke style. Has all the rest of the other features available on the color grouping except the advanced mode.

Group by Chart

The end result of using this feature: divides into subplots based on the value of the selected params. Why this is a grouping mechanism? It groups the metrics belonging to the same group into separate charts.

10.1.5 Aggregate metrics

The metrics aggregation helps to quickly see the trends of each individual group of metrics. See more about [metrics grouping](#).

There are two aspects of aggregation you can control:

- the trend-line
- the area that the group of metrics take

The trend-line calculation methods:

- Mean
- Median
- Min
- Max

The area calculation methods:

- None (*when you'd like to remove the area*)
- Min/Max
- Mean +/- Standard Deviation

- Mean +/- Standard Error
- Confidence Interval (95%)

Pls see the screenshot:

10.1.6 Axes properties

Axes properties section is for aligning metrics by time, epoch or another metric and for setting axes range manually.

Alignment:

Following types of metrics alignment are available: *Step*, *Epoch*, *Relative Time*, *Absolute Time* and *Custom Metric*. By default, metrics aligned by *Step*.

Step

By setting metrics alignment to Step, x-axis values will represent the steps of tracked metrics.

Epoch

By setting metrics alignment to Epoch, x-axis values will represent the epochs of tracked metrics.

Relative Time

By setting metrics alignment to Relative Time, x-axis values will represent by HH:mm:ss, duration of tracking process.

Absolute Time

By setting metrics alignment to Absolute Time, x-axis values will represent by date HH:mm:ss D MMM, YY, since the start date of the first run until the last run.

Custom Metric

By setting metrics alignment to Custom Metric, x-axis values will represent selected metric values, you can detect correlations between queried metrics and selected metric.

Set axes range:

To fix an axis range across all the charts, set the corresponding axis minimum and maximum bounds in the form.

10.1.7 Axes Scale

Axes Scale section gives ability to display axes scale's [linear](#) or [logarithmic](#).

By default, axes scale's are [Linear](#).

Linear Scale

X-axis scale: Linear, Y-axis scale: Log

X-axis scale: Log, Y-axis scale: Linear

Log Scale

10.1.8 Chart Smoothing

While smoothing the chart, the data points are modified so individual points higher than the adjacent points (presumably because of noise) are reduced, and points that are lower than the adjacent points are increased leading to a smoother signal. You can select curve interpolation methods: Linear or Cubic. By default, chart smoothing in *Exponential moving average* mode and curve interpolation method is Linear.

Exponential moving average

An **exponential moving average**, also known as an exponentially weighted moving average (EWMA), is a first-order infinite impulse response filter that applies weighting factors which decrease exponentially.

Centered moving average

When you center the moving averages, the data points are placed at the center of the range rather than the end of it. This is done to position the moving average values at their central positions in time.

10.1.9 Ignore outliers

An outlier is an observation that lies an abnormal distance from other values in a random sample from a population. Examination of the data for unusual observations that are far removed from the mass of data. These points are often referred to as outliers.

Excluding outliers can cause your results to become statistically significant. By default, outliers are ignored.

10.1.10 Highlight Modes

Highlighting functionality is useful for filtering metrics and highlight only hovered metric. Following types of highlighting mode are available: *Highlight Off*, *Highlight Metric on Hover*, and *Highlight Run on Hover*. By default, highlighting mode is the *Highlight Run on Hover*.

Highlight Off

By setting Highlight mode Off, there is no highlighting functionality on hover.

Highlight Metric on Hover

By setting Highlight mode Metric on Hover, mouse point closest metric highlights and other metrics displays with opacity.

Highlight Run on Hover

By setting Highlight mode Run on Hover, mouse point closest metric highlights and highlighted metric corresponding run also highlights other metrics displays with opacity.

10.1.11 Set tooltip parameters

You can select tooltip parameters to show params and values in tooltip Params section. You can select hide or show button to display or hide tooltip on hover.

10.1.12 Apply zoom on charts

Zoom In

Zoom Out

10.1.13 Export chart as image

Metric explorer also, gives ability to export your chart as image. By clicking `export` button from control panel, will be opened chart preview modal. You can change exportable chart image width, single chart height, set image name and format.

Following formats of chart export are available: SVG, JPEG, PNG.

10.2 Images Explorer

10.2.1 Overview

Track intermediate images search easily by using select form functional compare them on the Images Explorer by using reach controls panel.

Features:

- *Easily query any image*

- *Group images by run parameters, step and index*
- *Use controls from right control panel to configure workspace*
 - *Image properties control*
 - *Images sorting control*
 - *Images group stacking control*
 - *Images tooltip params control*

10.2.2 Query any image

Use select form to easily query any image. There are two option to query images using dropdown, by using [Aim QL](#) language and advance mode for [Aim QL](#).

- Click on Images button
 - Select options you are want to use in query
 - Click on the Search button
-
- Click on pencil icon in the right side of select form to show input
 - Type advance [Aim QL](#) query
 - Click on the Search button

10.2.3 Group image by run parameters

Use select grouping dropdown which is located in the right top corner of the image explore page.

- Click on grouping button
- Select fields by which you want to groupe images

Grouping will be apply after each field selection also you can select grouping mode (Group or Reverse)

10.2.4 Image explorer right controls panel

Any change in controls will help to explore images better on the workspace

10.2.5 Images size manipulation control

- Click on image property button
- Select value from dropdown to align image. (by default dropdown value is Height). Use slider to configure value for scale relative to window size by default scale value is 15%.
 - By height
 - By width
 - Original size
- Use image rendering variation by default value of this control is Pixelated

10.2.6 Images sorting control

- Click on image sorting button
- Select fields for sorting images. Selection ordering is meaningful and data will be sorting by selection order. Bellow is visible Ordered By list where contains all selected fields from dropdown. You can remove any already selected field by clicking on x icon or change sorting direction by clicking radio button Asc or Desc. Default selected direction is Asc.
- For reset all existing sorting fields you can simply click on Reset Sorting button

10.2.7 Set tooltip parameters

You can select tooltip parameters to show params and values in tooltip Params section. You can select hide or show button to display or hide tooltip on hover.

10.3 Params Explorer

10.3.1 Overview

Params explorer helps you to represent high dimensional data as a multi-dimensional visualization. Features:

- *Easily query any metrics and params*
- *Group runs by color, stroke, or chart*
- *Make the crossings easier with curve interpolation*
- *Learn patterns and colorations easier colored by the last dimension with a color indicator*

Query any metrics and params

Select params and metrics from dropdown

Search runs with [Aim QL](#)

Grouping

Group by color, stroke, or chart with selected parameters

Curve interpolation

By clicking on the Curve interpolation button in the Controls panel, it's possible to make lines from straight to curve to show correlations between non-adjacent axes.

Color indicator

By clicking on the Color indicator button in the Controls panel, it's possible to turn on lines gradient coloring by the last dimension.

10.4 Scatters Explorer

Scatter explorer gives ability to visualize correlations between metric last value data with hyper-parameter.

It represents graph where corresponding values from a set of data are placed as points on a coordinate plane. A relationship between the points is sometimes shown to be positive, negative, strong, or weak.

Abilities provided by Scatter explorer

- *Easily align metric last value data with hyper parameter*
- *Group runs by color and chart*
- *Represent the points with trend line*
- *Export chart*

Select params and metrics from X and Y axes dropdowns to align metric last value data with hyper-parameter.

- X axis
- Y axis

Also, you can search runs with [Aim QL](#)

Easily group data by color and chart with selected parameters.

- By Color
- By Chart

A trend line is a straight line that best represents the points on a `scatter plot`. The trend line may go through some points but need not go through them all.

From trend line options popover you can change regression from `Linear` (by default) to `LOESS`(locally weighted smoothing), which creates a smooth line through a `scatter plot` to help you to see relationship between variables and foresee trends. Also, you can change the `bandwidth` with `slider`

Scatter explorer also, gives ability to `export` your chart as `image`.

By clicking `export` button from control panel, will be opened chart preview modal. You can change exportable chart `image width`, `single chart height`, set `image name` and `format`.

Following image formats are available export: `SVG`, `JPEG`, `PNG`.

11.1 Overview

The Aim Table helps to explore and compare training runs - their results, logged hyper-parameters, system info, etc. You can also use the table to hide, delete, archive and compare selected training runs in a few clicks.

11.2 Delete the selected runs

To delete a set of runs, select rows and click on the **Delete** button, which is located at the top of the table.

In-progress runs can't be deleted. Please wait until the run is finished or terminate it to be able to delete the run.

After clicking on the **Delete** button a popup appears where **finished** and **in-progress** runs are divided into separate tables.

To undo the selection, hover over the selected row and click on the **X** button (appears on hover).

11.3 Archive/Unarchive the selected runs

To archive/unarchive a set of runs, select rows and click on the **Archive/Unarchive** button, which is located at the top of the table.

After clicking on the **Archive/Unarchive** button a popup appears, where archived and unarchived runs are divided into separate tables.

To undo the selection, hover over the selected row and click on the **X** button (appears on hover).

11.4 Hide/Show the selected runs

To hide/show a set of runs, select rows and click on the Hide/Show Lines button, which is located at the top of the table.

BOOKMARKS

12.1 Overview

Use the Bookmarks to save the Aim Explorer state. This includes search query, aggregations and any other modifications applied to the explorer. The Bookmarks page is a list of *cards* to quickly access the explorer state with one click.

There are bookmark buttons available on all Explorer pages - on the top right . In order to create a bookmark, just press the bookmark button on your preferred Explorer page.

This will open a `create bookmark` form with fields for `title` and `description`.

The created bookmark will contain all the current configuration of the explorer.

12.2 The Bookmark Card

Each of the bookmark cards contains the following:

- Explorer Icon
- Title
- View Bookmark button
- Delete Bookmark button
- Search query
- Selected metrics and params

12.3 Delete Bookmark

Clicking the delete bookmark button will open a modal, where you can confirm or cancel the deletion.

13.1 Overview

Tags functionality intended to mark a runs. A tag can be attached to the runs to distribute by segments and then find it quickly.

13.2 Create tag

How to create tag? There are two options for creating a tag.

13.2.1 First option

- Go to the tags page by clicking on the Tags from the left sidebar.
- Click on the create tag button to open the create tag form modal. In this modal there is the form that has two input fields first one for tag name the second one for tag comment and also there are exist color pallets for selecting tag color.
- Type name for a tag. Name filed is mandatory and can't be empty for tag creation form and has maximum 50 symbol limit validation.
- Type comment for a tag. Comment field is optional for tag creation form and has max 200 symbol limit validation.
- Select color for a tag from the color pallet.
- Click to the create button for saving a the tag then. After successful saving should appear toaster approving the create on the right top corner of the window. Optional there are default selected colors for tag.

13.2.2 Second option

- Go to any explorer page (metrics, params, images, scatters).
- Click to one of the sequence unit to open popover where is exist tag section with attach button.
- Click on attach button to open the select tag popover where you will see all your previously created tags.
- Click create tag button and you will be redirected to the tags page than the actual first option.

13.3 Attach tag

- Go to any explorer page (metrics, params, images, scatters).
- Click to one of the sequence unit to open popover where is exist tag section with attach button.
- Click on attach button to open the select tag popover where will be visible all tags.
- Select a tag you want to attach to the sequence unit. You can select more then on tag for each point.

13.4 Update attached tags

How to update attached tags?

- Go to any explorer page (metrics, params, images, scatters).
- Select point which you want to attach tag and click on it to open popover where is the exist tag section. In tags section will be visible already attached tags.
- Click on attach button for adding new tag to open the select tag popover where will be visible all existing tags.
- Click on x icon in the right end of the each tag for removing the tag from the point.

13.5 Edit tag

- Go to the tags page
- Click to edit icon in the right side in the tag row
- Then should appear the edit modal. In the edit modal there is the form that has two input fields first one for tag name the second one for tag comment and also there are exist color pallets for selecting tag color. In this modal is possible to make changes for tag.
- Make changes you need.

- Then you have three possible actions close modal, save changes and reset changes. After closing the modal you will lose all changes, after clicking the reset button modal form fields will be reset to initial values and by clicking the save button you will save all changes for the tag. After successful saving should appear toaster approving the update on the right top corner of the window.

13.6 Delete tag

- Go to the tags page.
- Click on the trash icon in the right side in the tag row to open the delete modal.
- In the delete modal there is a tag name input field and a tag name label at the top of the tag name input. You need to type the tag name for approving you are want to delete that tag.
- Then you have two possible actions delete the tag or close the modal by canceling the delete operation. If you want to delete a tag please double-check the tag name and click to delete button. After tag deletion, there are no possibilities to recover it. Also if you are deleting the tag it will be removed from all relations too.

13.7 Used in overlay

In the tags page you can select tag by clicking on circle icon then will opened overlay in the right side of window. Here is visible that runs which are use the tag. By clicking a run hash you will be redirected to single run page.

MANAGE RUNS

14.1 Create runs

Run is the main object that tracks and stores ML training metadata(e.g. metrics or hyperparams).

When initializing a Run object, Aim creates a `.aim` repository at the specified path. The tracked data is stored in the `.aim` repo. If the repo path is not specified, the run data is stored in the current working directory.

Use the following Run arguments to:

- `repo`: define where to store the data
- `experiment`: define experiment name to group related runs together
- `system_tracking_interval`: Enable system resource usage tracking (CPU, GPU, memory, etc..). By default enabled. Set to `None` to disable
- `capture_terminal_logs`: Enable process output capturing. By default enabled. Set to `False` to disable.

```
from aim import Run

my_run = Run(
    repo='/repo/path/to/store/runs',
    experiment='experiment_name'
)
```

Run class full `spec`.

Additionally, Aim SDK also gives a flexibility to:

- Use multiple Runs in one training script to store multiple runs at once. Usually handy when doing hyperparameter search.
- Use integrations to automate tracking

14.2 Continue runs

Each Run object has a hash associated with it which could be looked up at `aim runs ls` (check out the [Aim CLI](#) here). Specify the run hash when initializing a Run object to continue tracking.

```
from aim import Run

run = Run(run_hash='run_hash')
```

14.3 Delete runs

There are cases when Run data is not needed. Examples of such cases are, failed training runs or simple disk space cleanup. Aim provides SDK and CLI interfaces to delete Runs.

To remove Runs via the SDK:

```
from aim import Repo

repo = Repo.from_path('aim_repo_path')
repo.delete_run('run_hash')
repo.delete_runs(['run_hash_1', 'run_hash_2'])
```

Repo class full [spec](#).

To remove Runs using command line:

```
aim runs rm run_hash_1 run_hash_2 run_hash_3
```

14.3.1 Cleanup (prune) run params and metric names/contexts

Due to the Aim storage structure, after runs are deleted their params and metric names/contexts are still available for autocomplete in Aim UI explorers.

In order to clean up those dangling params and metric properties using command line:

```
aim storage prune
```

14.4 Upload runs

Aim provides a way to create backup snapshots of a `.aim` repository in AWS s3. The snapshot will be created in an s3 bucket. The bucket name is passed as an argument. A new s3 bucket will be created if it does not exist already.

Please note that `boto3` package and sufficient AWS permissions are required for this command.

The details of how to configure `boto3` credentials can be found [here](#).

To create a snapshot using command line:

```
aim runs upload bucket_name_1
```

More details on `aim runs` in CLI [reference](#).

CONFIGURE RUNS

You can configure Aim run instance to change the default repository location, retrieve a particular run instance, or store custom parameters for your run. You can configure your Run upon instantiation.

Here you can see all available options you can use to configure your Run instance.

- **run_hash**: Run's hash. Can be used to restore/retrieve a particular run from the repo. This is handy if you want to update something or add new metrics in your already existing run. If skipped, the hash will be generated automatically. Example covered in *Reusing Run in your Repo* section. Starting from Aim version 3.13, run_hash values restricted to the auto-generated values only. This is done to prevent multiple undesired cases, such as, creating a new Run instead of re-opening the existing one due to the typo, putting special symbols in Run.hash, etc. If the wrong run_hash is specified, `MissingRunError` will be raised.
- **repo**: Aim repository path or Repo object to which Run object is bound. If skipped, default Repo is used. More covered in *Defining custom location for your Repo* section.
- **read_only**: Run creation mode. Default is False, meaning Run object can be used to track metrics.
- **experiment**: Sets Run's experiment property. 'default' if not specified. Can be used later to query runs/sequences.
- **system_tracking_interval**: Sets the tracking interval in seconds for system usage metrics (CPU, Memory, etc.). Set to None to disable system metrics tracking.
- **log_system_params**: Enable/Disable logging of system params such as installed packages, git info, environment variables, etc. Checkout *Training Run Reproducibility* section for more info.

15.1 Reusing Run in your Repo

Sometimes you might want to add/update parameters and tracked values in your existing Run instance. This can be achieved through using run_hash parameter. Aim generates a unique id (uuid) for your Run instance only if run_hash parameter is not supplied. Example of re-using generated run.

```
from aim import Run

# You can also use ID shown in Aim Web after
# creating a run without supplying run_hash parameter
uid = '508c5b29-02c7-4875-a157-f099ea193bfa'
run = Run(run_hash=uid)
for i in range(100):
    run.track(i, step=i, name='test')
```

In the example above, no matter how often this piece of code will be executed, only a single run entry will be created or updated.

15.2 Defining custom location for your Repo

When you use aim, the library automatically creates `.aim` directory which we call “Repo” in your current directory. Aim stores data about your runs inside the Repo. Sometimes creating a Repo in your current directory is not convenient due to various reasons (like lack of permissions), and in such scenarios you can define your own path where a new `.aim` directory will be created or just reused.

```
from aim import Run

run = Run(repo='~/repo')

...
```

The code above will create a new dir entry `~/repo/.aim`

15.3 Organizing Runs in Experiments

Aim allows you to group runs under experiments, which can be useful for comparing runs intended to tackle a particular task.

```
from aim import Run

run = Run(experiment="fraud-detection")
```

15.4 Adding tags and params to Run

Aim lets you add tags and parameters on your Run object. Example of this would be if you want to tag your run with a specific version or keyword. Or you might need to log parameters, so you have some insight on what params have been used to process that particular Run.

Example of adding and removing tags

```
from aim import Run

run = Run()
run.add_tag('v1.0')
run.add_tag('some-awesome-tag')
```

Or you can modify tags on your existing run, but first you have to restore Run object by using its hash value

```
from aim import Run

uid = '508c5b29-02c7-4875-a157-f099ea193bfa'
run = Run(run_hash=uid)
run.remove_tag('some-awesome-tag')
run.add_tag('another-awesome-tag')
```

Almost same approach goes for parameters

```
from aim import Run
```

(continues on next page)

(continued from previous page)

```
run = Run()
var = 'some value'

run['var'] = var
```

Also, you can restore the Run object using its hash key and overwrite value of the previously stored parameter.

```
from aim import Run

uid = '508c5b29-02c7-4875-a157-f099ea193bfa'
run = Run(run_hash=uid)
var = 'another value'

run['var'] = var
```

15.5 Training Run Reproducibility

When running multiple training jobs it is crucial to understand the factors affecting the trained models performance. While the hyperparameters diff is an obvious place to look at, the training script environment itself can change the collected metadata in unexpected ways. It is important to be able to reproduce your runs' environment to presumably get the same results. Sometimes even minor version change in your script dependencies or a small tweak in the training code itself can affect training results. Thus, it's important to collect and log information such as package versions, environment variables, input arguments, etc. with each run.

Doing this manually requires a lot of code to be added to your training script. This is where Aim's logging of system parameters can come in handy!

15.5.1 What data is logged automatically?

Aim lets you enable system params logging for your Run which in result will log the following parameters

- Environment Variables
- Executables
- CLI arguments
- Installed packages and their versions
- Git information such as current branch, commit hash, author, etc. (if applicable)

15.5.2 How to enable system params automatic logging?

To enable logging of the parameters listed above, your Run instance must be supplied with `log_system_params=True` option, by default it is disabled!

```
run = Run(log_system_params=True)
```

In addition, these logged parameters can be used in the search box to filter runs based on the supplied parameters. Everything is searchable at Aim!

Here is an example of what you can do with it:

```
run.__system_params.git_info.branch == 'feature/testing'
```

15.6 Capturing Terminal logs

When it comes to automating training of multiple runs with job schedulers or workload managers on a cluster, it becomes hard to track the terminal logs of the runs. This includes the stdout, stderr from the training nodes too.

Inspecting these terminal logs is important, as it may convey crucial information about the training, for example:

- warnings, alerts and errors about the training process,
- debugging information,
- model training statistics, model architecture or summary.

Aim makes it is easier to analyze the terminal logs in one place by streaming them to Aim UI near-real-time.

Terminal logs are captured by default. In order to disable it, set the `capture_terminal_logs` argument to `False`, when initializing the Run:

```
aim_run = Run(capture_terminal_logs=False)
```

Check out how to view captured logs on Aim UI [here](#).

QUERY RUNS AND OBJECTS

Use Repo object to query and access saved Runs.

Initialize a Repo instance:

```
from aim import Repo

my_repo = Repo('/path/to/aim/repo')
```

Repo class full spec.

Query logged metrics and parameters:

```
query = "metric.name == 'loss'" # Example query

# Get collection of metrics
for run_metrics_collection in my_repo.query_metrics(query).iter_runs():
    for metric in run_metrics_collection:
        # Get run params
        params = metric.run[...]
        # Get metric values
        steps, metric_values = metric.values.sparse_numpy()
```

Besides querying Runs and metrics, you can also query logged Image objects:

```
query = "images.name == 'mnist_dataset' and images.context.subset == 'val'"

# Get collection of Image sequences
for image_seq in my_repo.query_images(query).iter():
    # Get first tracked batch of each sequence
    image_batch = image_seq.values.first_value()
    # Get Image metadata
    image_meta = map(Image.json, image_batch)
    # Convert to PILImage
    pil_images = map(Image.to_pil_image, image_batch)
```

Image class full spec.

See more advanced usage examples [here](#).

QUERY LANGUAGE BASICS

17.1 Introduction

Aim enables a powerful query language(Aim QL) to filter through all the stored metadata.

AimQL filters the tracked metadata using **python expression**. Think of it as a python if statement over everything you have tracked. Hence, nearly any python compatible expression is available with *some restrictions* in place.

The data is saved as different types of entities (e.g. `run`, `metric`). The search queries are written against these entities. When iterating over entities the python expression is evaluated in a Boolean context. When the value is “truthy”, then the current entity is yielded. Otherwise the entity is skipped over.

Currently, AimQL is only used for filtering data, and has no role in sorting or aggregating the data.

17.2 Searching

Let’s track several Runs via Aim SDK:

```
from aim import Run
# Initialize run_1
# Define its params and track loss metric within test and train contexts
run_1 = Run()
run_1['learning_rate'] = 0.001
run_1['batch_size'] = 32
for i in range(10):
    run_1.track(i, name='loss', context={ 'subset':'train' })
    run_1.track(i, name='loss', context={ 'subset':'test' })

# Initialize run_2
run_2 = Run()
run_2['learning_rate'] = 0.0007
run_2['batch_size'] = 64
for i in range(20):
    run_2.track(i, name='loss', context={ 'subset':'train' })
    run_2.track(i, name='loss', context={ 'subset':'test' })
    run_2.track(i/100, name='accuracy')

# Initialize run_3
run_3 = Run()
run_3['learning_rate'] = 0.005
```

(continues on next page)

(continued from previous page)

```
run_3['batch_size'] = 16
for i in range(30):
    run_3.track(i, name='loss', context={ 'subset':'train' })
    run_3.track(i, name='loss', context={ 'subset':'test' })
    run_3.track(i/100, name='accuracy')
```

Aim SDK will collect and store the above metadata in .aim repo.

Run	Parameters	Metrics
run_1 <hash=a32c910>		
run_2 <hash=a32c911>		
run_3 <hash=a32c912>		

When searching runs, use the run keyword which represents the Run object. It has the following properties:

Property	Description
name	Run name
hash	Run hash
experiment	Experiment name
tags	List of run tags
archived	True if run is archived, otherwise False
active	True if run is active(in progress), otherwise False
duration	Run duration in seconds
created_at	Run creation datetime
finalized_at	Run end datetime
metrics	Set of run metrics

Run parameters could be accessed both via chained properties and attributes.

The three following examples are functionally equal:

- `run.hparams.learning_rate == 32` (recommended)
- `run["hparams", "learning_rate"] == 32` (recommended)
- `run["hparams"]["learning_rate"] == 32`

AimQL has been designed to be highly performant. Only the params that are used in the query will be loaded into memory.

If you use the `['hparams']['learning_rate']` syntax Aim will load the whole dictionary into memory. The search performance will be impacted.

We recommend to use either `['hparams', 'learning_rate']` or `hparams.learning_rate` syntax which are equivalent to each other in terms of the performance.

Query examples:

1. Get runs where `learning_rate` is greater than 0.0001 and `batch_size` is greater than 32.

```
run.learning_rate > 0.0001 and run.batch_size > 32
```

Result:

Run	Parameters
run_2 <hash=a32c911>	

1. Get runs where learning_rate is either 0.0001 or 0.005.

```
run.learning_rate in [0.0001, 0.005]
```

Result:

Run	Parameters
run_1 <hash=a32c910>	
run_3 <hash=a32c912>	

1. Search runs based on metrics' last values.

metrics property takes 2 arguments: metric name and context (name is required, context is empty({}) by default).

- a. Get runs where accuracy, {} last value is greater than 0.25

```
run.metrics['accuracy'].last > 0.25
```

Result:

Run	Parameters	Metrics
run_3 <hash=a32c912>		

- b. Get runs where loss, {'subset': 'train'} last value is smaller than 25

```
run.metrics['accuracy', {'subset': 'train'}].last < 25
```

Result:

Run	Parameters	Metrics
run_1 <hash=a32c910>		
run_2 <hash=a32c911>		

1. Search runs where duration is smaller than 10s (to filter out failed runs for example) .. code-block:: python

```
run.duration < 10
```

This will return all the runs on our test set, but will have a use case on real life experiments to filter the runs that finished under 10 seconds. Then those can be archived/deleted as they've probably crashed or were stopped.

1. Search runs based on creation or end times

created_at and finalized_at properties are represented as python datetime objects. So any of these queries are a competent way to filter runs:

- a. Get runs for a specific date

```
run.created_at.strftime("%d/%m/%Y") == "24/12/2021"
```

b. Get runs after a specific date

```
run.finalized_at > datetime(2022, 2, 2)
```

c. Get runs of a specific month

```
run.created_at.month == 2 and run.created_at.year == 2022
```

17.3 Searching metrics and images

17.3.1 Searching metrics

When iterating over metrics, use the `metric` keyword which represents the tracked metric. While searching metrics, you can also refer to the related runs via the `run` keyword.

`metric` has the following default properties.

Property	Description
<code>name</code>	Metric name
<code>context</code>	Metric context dictionary
<code>last</code>	Last tracked value of the metric
<code>last_step</code>	Last tracked step number of the metric
<code>first_step</code>	First tracked step number of the metric

Query examples

1. Query metrics by name:

```
metric.name == "loss"
```

Result:

Metric	Related run
loss { "subset": "train" }	run_1 <hash=a32c910>
loss { "subset": "test" }	run_1 <hash=a32c910>
loss { "subset": "train" }	run_2 <hash=a32c911>
loss { "subset": "test" }	run_2 <hash=a32c911>
loss { "subset": "train" }	run_3 <hash=a32c912>
loss { "subset": "test" }	run_3 <hash=a32c912>

1. Query metrics by name and context

```
metric.name == "loss" and metric.context.subset == "train"
```

Result:

Metric	Related run
loss { "subset": "train" }	run_1 <hash=a32c910>
loss { "subset": "train" }	run_2 <hash=a32c911>
loss { "subset": "train" }	run_3 <hash=a32c912>

1. Query metrics by name and run parameters

```
metric.name == "loss" and run.learning_rate >= 0.001
```

Result:

Metric	Related run
loss { "subset": "train" }	run_1 <hash=a32c910>
loss { "subset": "test" }	run_1 <hash=a32c910>
loss { "subset": "train" }	run_3 <hash=a32c912>
loss { "subset": "test" }	run_3 <hash=a32c912>

1. Query metrics by name and last value

```
metric.name == "loss" and metric.last >= 15
```

Result:

Metric	Related run
loss { "subset": "train" }	run_2 <hash=a32c911>
loss { "subset": "test" }	run_2 <hash=a32c911>
loss { "subset": "train" }	run_3 <hash=a32c912>
loss { "subset": "test" }	run_3 <hash=a32c912>

17.3.2 Searching images

Images search works in the same way as metrics. When iterating over images, use the `images` keyword which represents the tracked images sequence. While searching images, you can also refer to the related runs via the `run` keyword.

`images` keyword has the following default properties.

Property	Description
<code>name</code>	Image sequence name
<code>context</code>	Image sequence context dictionary

Query examples:

- `images.name == "generated" and run.learning_rate >= 0.001`
- `images.name == "generated" and images.context.ema == 0`

17.4 Security restrictions

AimQL expression is evaluated with `RestrictedPython`.

RestrictedPython is a tool that helps to define a subset of the Python language which allows to provide a program input into a trusted environment.

We have followed these [restrictions](#) to avoid security risks such as executing a non-safe function via AimQL.

TRACK EXPERIMENTS WITH AIM REMOTE SERVER

18.1 Overview

Aim remote tracking server allows running experiments in a multi-host environment and collect tracked data in a centralized location. It provides SDK for client-server communications and utilized `gRPC` protocol as its core transport layer.

In this guide we will show you how to setup Aim remote tracking server and how to integrate it in client-side code.

18.2 Prerequisites

Remote tracking server assumes multi-host environments used to run multiple training experiments. The machine running the server have to accept incoming TCP traffic on a dedicated port (default is 53800).

18.3 Scaling

Aim remote tracking server can be set up to use multiple workers to scale the bandwidth of receiving connections. But it comes with some caveats. When providing `--workers=n` argument to `aim server` command, `aim` uses `n + 1` sequential ports starting from port number specified. Thus, all `p, ..., p + n` ports are required to be exposed to the client. No changes are needed on client side.

Note:

The default number of workers is 1. No additional ports are required to be opened in this case.

18.4 Server-side setup

1. Make sure `aim 3.4.0` or upper installed:

```
$ pip install "aim>=3.4.0"
```

2. Initialize `aim` repository (optional):

```
$ aim init
```

1. Run `aim` server with dedicated `aim` repository:

```
$ aim server --repo <REPO_PATH>
```

You will see the following output:

```
> Server is mounted on 0.0.0.0:53800
> Press Ctrl+C to exit
```

The server is up and ready to accept tracked data.

1. Run aim UI

```
$ aim up --repo <REPO_PATH>
```

18.5 Client-side setup

With the current architecture there is almost no change in aim SDK usage. The only difference from tracking locally is that you have to provide the remote tracking URL instead of local aim repo path. The following code shows how to create Run with remote tracking URL and how to use it.

```
from aim import Run

aim_run = Run(repo='aim://172.3.66.145:53800') # replace example IP with your tracking_
↪server IP/hostname

# Log run parameters
aim_run['params'] = {
    'learning_rate': 0.001,
    'batch_size': 32,
}
...
```

You are now ready to use aim_run object to track your experiment results. Below is the full example using pytorch + aim remote tracking on MNIST dataset.

```
from aim import Run

import torch
import torch.nn as nn
import torchvision
import torchvision.transforms as transforms

# Initialize a new Run with remote tracking URL
aim_run = Run(repo='aim://172.3.66.145:53800') # replace example IP with your tracking_
↪server IP/hostname

# Device configuration
device = torch.device('cpu')

# Hyper parameters
num_epochs = 5
num_classes = 10
```

(continues on next page)

(continued from previous page)

```

batch_size = 16
learning_rate = 0.01

# aim - Track hyper parameters
aim_run['hparams'] = {
    'num_epochs': num_epochs,
    'num_classes': num_classes,
    'batch_size': batch_size,
    'learning_rate': learning_rate,
}

# MNIST dataset
train_dataset = torchvision.datasets.MNIST(root='./data/',
                                           train=True,
                                           transform=transforms.ToTensor(),
                                           download=True)

test_dataset = torchvision.datasets.MNIST(root='./data/',
                                           train=False,
                                           transform=transforms.ToTensor())

# Data loader
train_loader = torch.utils.data.DataLoader(dataset=train_dataset,
                                           batch_size=batch_size,
                                           shuffle=True)

test_loader = torch.utils.data.DataLoader(dataset=test_dataset,
                                           batch_size=batch_size,
                                           shuffle=False)

# Convolutional neural network (two convolutional layers)
class ConvNet(nn.Module):
    def __init__(self, num_classes=10):
        super(ConvNet, self).__init__()
        self.layer1 = nn.Sequential(
            nn.Conv2d(1, 16, kernel_size=5, stride=1, padding=2),
            nn.BatchNorm2d(16),
            nn.ReLU(),
            nn.MaxPool2d(kernel_size=2, stride=2))
        self.layer2 = nn.Sequential(
            nn.Conv2d(16, 32, kernel_size=5, stride=1, padding=2),
            nn.BatchNorm2d(32),
            nn.ReLU(),
            nn.MaxPool2d(kernel_size=2, stride=2))
        self.fc = nn.Linear(7 * 7 * 32, num_classes)

    def forward(self, x):
        out = self.layer1(x)
        out = self.layer2(out)
        out = out.reshape(out.size(0), -1)
        out = self.fc(out)

```

(continues on next page)

(continued from previous page)

```

    return out

model = ConvNet(num_classes).to(device)

# Loss and optimizer
criterion = nn.CrossEntropyLoss()
optimizer = torch.optim.Adam(model.parameters(), lr=learning_rate)

# Train the model
total_step = len(train_loader)
for epoch in range(num_epochs):
    for i, (images, labels) in enumerate(train_loader):
        images = images.to(device)
        labels = labels.to(device)

        # Forward pass
        outputs = model(images)
        loss = criterion(outputs, labels)

        # Backward and optimize
        optimizer.zero_grad()
        loss.backward()
        optimizer.step()

    if i % 30 == 0:
        print('Epoch [{}/{}], Step [{}/{}], '
              'Loss: {:.4f}'.format(epoch + 1, num_epochs, i + 1,
                                      total_step, loss.item()))

    # aim - Track model loss function
    aim_run.track(loss.item(), name='loss', epoch=epoch,
                  context={'subset': 'train'})

    correct = 0
    total = 0
    _, predicted = torch.max(outputs.data, 1)
    total += labels.size(0)
    correct += (predicted == labels).sum().item()
    acc = 100 * correct / total

    # aim - Track metrics
    aim_run.track(acc, name='accuracy', epoch=epoch, context={'subset': 'train'})

    if i % 300 == 0:
        aim_run.track(loss.item(), name='loss', epoch=epoch, context={'subset':
        ↪ 'val'})
        aim_run.track(acc, name='accuracy', epoch=epoch, context={'subset': 'val
        ↪ '})

# Test the model

```

(continues on next page)

(continued from previous page)

```

model.eval()
with torch.no_grad():
    correct = 0
    total = 0
    for images, labels in test_loader:
        images = images.to(device)
        labels = labels.to(device)
        outputs = model(images)
        _, predicted = torch.max(outputs.data, 1)
        total += labels.size(0)
        correct += (predicted == labels).sum().item()

    print('Test Accuracy: {} %'.format(100 * correct / total))

```

18.6 SSL support

Aim Remote Tracking server can be set up to use SSL certificates to run on secure mode. To enable secure mode for server provide `--ssl-keyfile` and `--ssl-certfile` arguments to `aim servercommand`, where `--ssl-keyfile` is the path to the private key file for the certificate and `--ssl-certfile` is the path to the signed certificate (check out the [Aim CLI](#) here).

In order to establish a secure connection with the server, the client has to be configured accordingly. Please set `__AIM_CLIENT_SSL_CERTIFICATES_FILE__` environment variable to the file, where PEM encoded root certificates are located, e.g.

```
export __AIM_CLIENT_SSL_CERTIFICATES_FILE__=/path/of/the/certs/file
```

18.7 Message size limitations

Aim Remote Tracking server uses gRPC as a transport layer. gRPC imposes message size limits on sending/receiving messages from/to server. Aim is configured to use 16MB message size limit by default. If you want to specify a different limit, use `__AIM_RT_MAX_MESSAGE_SIZE__` environment variable.

```

# max message size 128MB
export __AIM_RT_MAX_MESSAGE_SIZE__=134217728

```

****Note:**** The message max size should be the same for both Aim Remote Tracking server and client code.

18.8 Conclusion

As you can see, aim remote tracking server allows running experiments on multiple hosts with simple setup and minimal changes to your training code.

NOTIFY ON FAILED/STUCK RUNS

Many things can cause the training process to fail or get stuck:

- hardware failure
- power outage
- programmatic error

At first, it may seem a relatively easy, however checking that everything is OK with the training run is proven to be a non-trivial task.

In some cases the training process may not even have a chance to notify about the failure. The process could also be stuck due to

- network issues
- filesystem I/O issues
- etc

With this in mind, the following definition of failed process is used:

The Run is considered as failed if it hasn't reported any progress in a predefined time-interval.

Once the progress is not reported for the given period of time, a notification will be sent to the enabled channels.

See how to set up the notification service and add notification channels [here](#).

19.1 Run progress reporting SDK

Aim SDK provides interface for reporting Run progress. The Run class has now two methods: `report_progress` and `report_successful_finish` to report progress and successful finish respectively. Here's a small code snippet showing how `Run.report_progress()` method can be integrated in the training loop.

```
from aim import Run

# prep dataset and model
...
aim_run = Run()

for epoch in range(num_epochs):
    for i, (images, labels) in enumerate(train_loader):
        images = images.to(device)
        labels = labels.to(device)
```

(continues on next page)

(continued from previous page)

```
# Forward pass
outputs = model(images)
loss = criterion(outputs, labels)

# Report progress assuming each iteration should take less than 3 sec.
aim_run.report_progress(expect_next_in=3)

# Backward and optimize
optimizer.zero_grad()
loss.backward()
optimizer.step()

if i % 30 == 0:
    # Track data with aim
    aim_run.track(loss.item(), name='loss', epoch=epoch)

...
# Training is done, report success
aim_run.report_successful_finish()
```

`report_progress()` takes an ETA (in seconds) of the next anticipated progress report call.

Detailed description of interfaces is available in aim [SDK reference](#).

Note: Additional grace period of 100s is enabled to compensate for possible hardware (e.g. filesystem) latency.

LOG MESSAGES DURING TRAINING PROCESS

During the training many key events may happen. Use Aim logging API to log any text message.

Additionally, [set up the notification service](#) to automatically send notifications, when a new message is logged.

20.1 The logging API

Aim SDK provides an interface to log any string message in the “.aim” repo.

- **Run.log_debug(msg)** - aim.Run method to log a message with level DEBUG.
- **Run.log_info(msg)** - aim.Run method to log a message with level INFO.
- **Run.log_warning(msg)** - aim.Run method to log a message with level WARNING.
- **Run.log_error(msg)** - aim.Run method to log a message with level ERROR.

20.2 Notifying on logged messages

Aim notification service constantly checks for new logged messages and sends notifications to the enabled channels. It is also possible to [configure](#) the level of notifications to be sent.

See how to configure notifiers [here](#).

SET UP THE NOTIFICATION SERVICE

Being notified about key events of the training process may save hours of compute resources and researcher's time. Aim notification service sends notifications to the enabled channels (slack, workplace), when a new message is logged or a stuck run is detected.

21.1 Quick start

These are the required steps in order to set up the notifications service.

1. Configure notifiers. The notifier configuration is described in detail in *Configuring Notifiers* section.
2. Start notifier service:

```
aim-watcher --repo <AIM_REPO_PATH> start
```

Aim notifications service is comprised of the following components:

- Aim watcher service. A service responsible for sending notifications.
- Notifier configuration CLI. Command line utility for configuring where and how you want to receive notifications about training run status.

Below are the components described in more detail.

21.2 Configuring notifiers

Use `aim-watcher CLI` to interactively choose and setup the notifiers. Currently, following types of notifiers are supported:

- *Slack*: to receive notifications as a message on a slack channel.
- *Workplace*: to receive notifications as a post in a workplace group.
- *Logger*: to echo notification text to watcher standard output.

The base command is:

```
aim-watcher notifiers add
```

Once notifier type is selected, the watcher will ask to fill-in the required inputs for that particular notifier.

`aim-watcher` allows configuring multiple notifiers in case you'd like to receive them in multiple channels at the same time. Once notifier configuration added, `aim-watcher` should be re-run, so new configuration can take effect.

Next sections will describe the setup flow for each of them.

21.2.1 Configuring notifier for slack

In order to execute the steps described below you must have admin permissions for Slack workspace.

Here are the steps to configure Aim notifier for Slack.

1. Go to Slack API web [page](#) and push “Create an app” button.
2. Select “From scratch” and follow the steps.
3. Make sure Incoming Webhooks are active. This is the Slack API Aim will use to send notifications.
4. Create a new Webhook URL and copy to clipboard.
5. Run `aim-watcher notifiers add slack` and once prompted for Webhook url, paste the one copied in step 4.
6. *Optional* Change the training run failure message template.
7. Confirm setup.

21.2.2 Configuring notifier for workplace

In order to execute steps described below, please make sure you have access to Workplace Admin panel

Here are the steps to configure Aim notifier for Workplace.

#. Go to “Your workplace” Admin Panel #.

Select “Integrations” tab and press “Create custom integration”

1. Follow the setup steps.
2. In the “Permissions” tab make sure “Manage group content” checkbox checked.
3. Select the group(s) where you want to receive notifications.
4. In the “Details” tab press “Create access token” and copy the generated access token.
5. Run `aim-watcher notifiers add workplace` and provide group ID and copied access token.
6. *Optional* Change the training run failure message template.
7. Confirm setup.

21.2.3 Enabling/disabling notifier configuration

Aim watcher CLI allows to temporarily disable or completely remove notifiers. Each notifier configuration has unique auto-generated ID which can be used to manipulate the notifier. In order to list available notifiers for the aim Repo:

```
aim-watcher --repo . notifiers list
```

> NOTIFIER ID	TYPE	STATUS
> 88bbacb3-159d-4ee2-8f33-180add162f1	slack	enabled
> 75fe529d-c060-4816-9cbb-246770cb8592	slack	disabled
> 248821e4-2bf3-4aa9-857e-194bf3ab2f8a	logger	enabled

You can disable notifier by given ID:

```
aim-watcher --repo . notifiers disable 88bbacb3-159d-4ee2-8f33-180add162f1
```

and enable it back:

```
aim-watcher --repo . notifiers enable 88bbacb3-159d-4ee2-8f33-180add162f1
```

In order to remove notifier configuration completely:

```
aim-watcher --repo . notifiers remove 88bbacb3-159d-4ee2-8f33-180add162f1
```

Complete reference for aim-watcher CLI is available [here](#).

21.2.4 Configuring notification levels

The notification levels to notify on logged messages can be configured. Available options are DEBUG, INFO, WARNING, ERROR. Run `aim-watcher notifiers set-log-level [LVL]` command to set the level.

Note: WARNING is set as a default level.

DEFINE CUSTOM CALLBACKS

Many things can go wrong during ML training (incorrect driver versions, plateaued metrics, etc) that could result in wasted GPUs and time. The Aim callbacks API helps to define custom callbacks to be executed at any point during ML training - a programmable way to guard ML training from wasting resources.

Callbacks can actually encompass any programmable functionality, such as [logging messages and sending notifications](#), or killing the training process when the given condition is met.

22.1 Callbacks

Terms:

- **callback** - python function that implements a custom logic to be executed at a certain point during the training.
- **callbacks class** - python class to group callback functions. Can be used to share state between different callbacks (example below).
- **event** - represents an event to be bound to the training.

The callbacks API:

- **TrainingFlow** - defines the training flow/events.
- **events.on.*** - decorators to define when the callback function is executed.

The list of currently available training events:

- **events.on.training_started** - called after the training start.
- **events.on.training_successfully_finished** - called after the training is successfully finished, meaning no unexpected exceptions are raised, even a manual keyboard interruption (ctrl+C). Please note that programmatic early stopping is considered as a successful finish.
- **events.on.training_metrics_collected** - called after the training metrics are calculated and ready to be logged. Typically called at each N batches.
- **events.on.validation_metrics_collected** - called after the validation metrics are calculated and ready to be logged. Mostly called only once, after the validation loop is finished.
- **events.on.init** - automatically called after the callbacks class initialization and before all the other events. Must not be called manually. Typical use-case can be initializing a shared state for callback functions (example below).

22.2 Example

The below example demonstrates how to implement custom callbacks to check and notify, when:

- wrong driver versions are installed.
- gnrm metrics explode.
- model starts to overfit.

22.2.1 Defining the callbacks

```
from aim.sdk.callbacks import events

class MyCallbacks:
    @events.on.init # Called when initializing the TrainingFlow object
    def init_gnorm_accumulators(self, **kwargs):
        # Initialize a state to collect gnorm values over training process
        self.gnorm_sum = 0
        self.gnorm_len = 0

    @events.on.init
    def init_ppl_accumulators(self, **kwargs):
        # Initialize a state to collect ppl values over training process
        self.ppl_sum = 0
        self.ppl_len = 0

    @events.on.init
    def init_metrics_accumulators(self, **kwargs):
        # Collect only the last 100 appended values
        self.last_train_metrics = deque(maxlen=100)

    # NOTE: all the above methods can be merged into one,
    #         but are separated for readability reasons

    @events.on.training_started
    def check_cuda_version(self, run: aim.Run, **kwargs):
        if run['__system_params', 'cuda_version'] != '11.6':
            run.log_warning("Wrong CUDA version is installed!")

    @events.on.training_metrics_collected
    def check_gnorm_and_notify(
        self,
        metrics: Dict[str, Any],
        step: int,
        # always denotes the number of *training* steps
        # `1 step per 4 batches` can be in case of gradient accumulation
        epoch: int,
        run: aim.Run,
        **kwargs
    ):
        current = metrics['gnorm'] # notice that it's the last one
        # thus we need to use self.* to collect gnorm values
```

(continues on next page)

(continued from previous page)

```

self.gnorm_sum += current
self.gnorm_len += 1
mean = self.gnorm_sum / self.gnorm_len

if current > 1.15 * mean:
    run.log_warning(f'gnorms have exploded. mean: {mean}, '
                   'step {step}, epoch {epoch} ...')

@events.on.training_metrics_collected
def check_ppl_and_notify(
    self,
    metrics: Dict[str, Any],
    step: int,
    epoch: int,
    run: aim.Run,
    **kwargs
):
    current = metrics['ppl'] # notice that it's the last one
    # thus we need to use self.* to collect ppl values
    self.ppl_sum += current
    self.ppl_len += 1
    mean = self.ppl_sum / self.ppl_len

    if current > 1.15 * mean:
        run.log_warning(f'ppl have exploded. mean: {mean}, '
                       'step {step}, epoch {epoch} ...')

@events.on.training_metrics_collected
def store_last_train_metrics(
    self,
    metrics: Dict[str, Any],
    step: int,
    epoch: int,
    **kwargs,
):
    self.last_train_metrics.append(metrics)

@events.on.validation_metrics_collected
def check_overfitting(
    self,
    metrics: Dict[str, Any],
    epoch: int = None,
    run: aim.Run,
    **kwargs,
):
    mean_train_ppl = sum(
        metrics['ppl'] for metrics
        in self.last_train_metrics
    ) / len(self.last_train_metrics)

    if mean_train_ppl > 1.15 * metrics['ppl']:
        run.log_warning(f'I think we are overfitting on epoch={epoch}')

```

22.2.2 Registering the callbacks

```
from aim import TrainingFlow, Run

aim_run = Run()

training_flow = TrainingFlow(run=aim_run, callbacks=[MyCallbacks()])
# or
training_flow = TrainingFlow(run=aim_run)
training_flow.register(MyCallbacks())
```


HOST AIM ON KUBERNETES (K8S)

Since Aim can run as a local server through FastAPI, it can be deployed to a K8S cluster! Hosting Aim on K8S comes with several advantages:

- multiple users of your organization can access Aim in a single spot, which removes the need for ML practitioners to run Aim themselves
- Aim runs can be centralized on a remote volume, which provides additional support and encouragement for remote model training and monitoring
- a deployment to K8S abstracts away the Aim CLI, which empowers users to focus on the value provided by Aim (visualizations/applications vs. CLI up and repo understanding)

The following sections illustrate how to deploy and serve Aim on K8S. The sections assume:

- access to a cloud provider, such as GCP, AWS, or Azure
- a repository that can host Dockerfiles, such as Google Artifact Registry or Dockerhub
- ability/permissions to provision a ReadWriteMany volume, or bind an existing one to a K8S deployment

23.1 Dockerfile

The following Dockerfile image should suffice for getting Aim running in a container:

```
# python3.7 should be sufficient to run Aim
FROM python:3.7

# install the `aim` package on the latest version
RUN pip install --upgrade aim

# make a directory where the Aim repo will be initialized, `/aim`
RUN mkdir /aim


ENTRYPOINT ["/bin/sh", "-c"]

# have to run `aim init` in the directory that stores aim data for
# otherwise `aim up` will prompt for confirmation to create the directory itself.
# We run aim listening on 0.0.0.0 to expose all ports. Also, we run
# using `--dev` to print verbose logs. Port 43800 is the default port of
# `aim up` but explicit is better than implicit.
CMD ["echo \"N\" | aim init --repo /aim && aim up --host 0.0.0.0 --port 43800 --workers_
↪2 --repo /aim"]
```

Assuming you store the above in your current directory, the container can be built using `docker build . -t my-aim-container:1` and pushed to your repository with `docker push my-docker-repository.dev/deployments/aim:1`.

23.2 Volume

The core advantage of using a K8S volume to store Aim runs is that other K8S deployments can mount the same volume and store their runs on it! This way, the core Aim K8S deployment can read the new runs and display them to users who want to visualize their results. For example, one can have a deployment that performs model training and records Aim runs on the same volume that is mounted to the Aim deployment! This model is illustrated by the following diagram:



Generally, volumes that support the `ReadWriteMany` property are manually provisioned, such as Filestore instances on Google Cloud or, generally, GlusterFS volumes. Once a disk is provisioned, it can be bound to a persistent volume via an IP. Assuming you can provision a disk like this on your cloud provider and obtain an IP, we can create a volume representation, along with a claim for it. The persistent volume (`aim-pv.yaml`) can be formulated as:

```

apiVersion: v1
kind: PersistentVolume
metadata:
  name: aim-runs
spec:
  capacity:
    storage: 1Ti # or whatever size disk you provisioned
  accessModes:
    - ReadWriteMany
  nfs:
    path: /aim
    server: 123.12.123.12 # add your own IP here
  
```

The persistent volume claim (`aim-pvc.yaml`) is:

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: aim-runs-claim
spec:
  accessModes:
  
```

(continues on next page)

(continued from previous page)

```

- ReadWriteMany
storageClassName: "" # if you have a custom storage class, use it! Otherwise, it's
↪ `default`
volumeName: aim-runs
resources:
  requests:
    storage: 1Ti

```

These can be provisioned via:

```

> kubectl apply -f aim-pv.yaml
> kubectl apply -f aim-pvc.yaml

```

Once the volume is provisioned, we can mount it to our deployments!

23.3 Deployment

The main Aim deployment will have a single container that runs Aim. This deployment will mount the volume that was provisioned previously, and the main Aim repository will be initialized at the path the volume is mounted to. For example, if the volume is mounted to /aim, then the deployment will initialize and read Aim runs from that path. The K8S deployment is:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app: my-aim-deployment
  name: my-aim-deployment
  namespace: default
spec:
  selector:
    matchLabels:
      app: my-aim-deployment
  strategy:
    rollingUpdate:
      maxSurge: 25%
      maxUnavailable: 25%
    type: RollingUpdate
  template:
    metadata:
      labels:
        app: my-aim-deployment
    spec:
      containers:
        image: my-docker-repository.dev/deployments/aim:1
        name: my-aim-deployment
        ports:
          - containerPort: 43800
            protocol: TCP
        resources:
          limits:

```

(continues on next page)

(continued from previous page)

```
    cpu: "1"
    memory: 2Gi
  requests:
    cpu: 100m
    memory: 1Gi
  volumeMounts:
    - mountPath: /aim
      name: aim-runs
  volumes:
    - name: aim-runs
      persistentVolumeClaim:
        claimName: aim-runs-claim
```

This K8S deployment:

- defines a pod with a single replica that runs the Aim server defined by the Dockerfile
- mounts the persistent volume `aim-run` through the `aim-run-claim` persistent volume claim
- the Dockerfile initializes the `/aim` directory as the Aim repo. Note that the Dockerfile already passes `N` to the confirmation prompt in case the repo is already initialized (this will be the case after the initial deployment creation, since the repo has to be initialized only once, but it's nice to avoid manual work)
- starts up the Aim server on port 43800, which reads all the runs stored at `/aim`

23.4 Service

Now that a deployment is deployed, the Aim server can be exposed through a K8S service! Depending on your cluster setup, you have several options for exposing the deployment. One option is to run:

```
> kubectl expose deployment my-aim-deployment --type=LoadBalancer --name=my-aim-service
```

Another alternative is to create the service definition yourself, and apply it. The definition (`aim-svc.yaml`) can be:


```
apiVersion: v1
kind: Service
metadata:
  name: my-aim-service
spec:
  selector:
    app: my-aim-deployment
  ports:
    - protocol: TCP
      port: 80
      targetPort: 43800
```

The service definition can be applied via:

```
> kubectl apply -f aim-svc.yaml
```

23.5 Conclusion

That's it! Now you have the following structure serving your users' Aim runs:



Assuming your users can submit a model training run to *some* pod/deployment that runs model training and has the right aim code to record a run at path `/aim`, your Aim deployment will be able to display the run the next time it performs a live update!

RUN AIM UI ON JUPYTER NOTEBOOK

Start your notebook with the following code to install Aim:

```
!pip install aim
```

Next, initialize a new run and save some hyperparameters:

```
from aim import Run

run = Run()

run['hparams'] = {
    'learning_rate': 0.001,
    'batch_size': 32,
}
```

Note: Do not forget to call `run.finalize()` once the training is over.

After tracking runs with `Run`, run the following commands in the notebook to run the Aim UI:

1. Load Aim extension for notebooks:

```
%load_ext aim
```

1. Run `%aim up` to open Aim UI in the notebook:

```
%aim up
```


RUN AIM UI ON SAGEMAKER NOTEBOOK INSTANCE

In this guide you will learn how to run Aim UI on your Sagemaker Jupyter instance

Aim can be installed and used on Jupyter notebooks. With SageMaker notebook instance, there are difficulties with establishing POST requests which prevents some Aim functionalities to work. So we have built a script to update the proxying package on SageMaker notebook instance. In order to work with Aim UI, please follow the steps below:

- Create a lifecycle configuration on your SageMaker Service.
- Copy and paste the [script](#) in your configuration's Start Notebook phase.

For more information how to create a lifecycle configuration on AWS SageMaker Service, please go through the [AWS SageMaker Documentation](#).

After creating a lifecycle configuration, attach it to your SageMaker instance and restart the instance.

It is possible to run this script manually from the terminal of your jupyter instance as well.

Note: Why this script is used for?

The script above will update the *jupyter-server-proxy* package on your notebook instance. To access a port on your instance (i.e. open the Aim UI) API requests and static files need to go through the proxy, which is created by the package (see the script). It will replace the existing proxy and make sure Aim and other such requests go through. For more information please go through the <https://github.com/jupyterhub/jupyter-server-proxy/pull/328#issue-1145874348>.

Once your notebook instance is successfully restarted you will be able to use Aim UI on the instance.

Next, install aim.

```
$ pip install aim
```

Initialize a new run and save some hyperparameters.

```
from aim import Run

run = Run()

run['hparams'] = {
    'learning_rate': 0.001,
    'batch_size': 32,
}
```

25.1 Using Terminal

In order to run Aim UI from a terminal, execute the following command:

```
$ aim up --base-path=/proxy/absolute/<your-port>/aim-sage
```

After running this command you will be able to open `<sagemaker_instance>/proxy/absolute/<your-port>/aim-sage/` in your browser. The default port is 43800.

It is possible to set `__AIM_PROXY_URL__` env variable, and `aim up` command will print out the generated url for Aim UI.

To find your proxy url, just copy your SageMaker URL and remove `/lab` postfix.

25.2 Using Notebook Extension

1. Load Aim extension for notebooks:

```
%load_ext aim
```

1. Run `%aim up` command to open Aim UI in the notebook:

```
%aim up --proxy-url=https://<instance>.notebook.<region>.sagemaker.aws
```

Will load the Aim UI in the notebook cell.

Will print out the proxy url under the loaded Aim UI. Use that URL to open Aim UI on your browser.

The default port is 43801 for notebook extension to prevent confusions. The `--port=<your-port>` argument is supported as well.

Note: In notebook extension, the only mandatory argument is `--proxy-url` when using it on SageMaker.

INTEGRATION GUIDES

Aim integrates seamlessly with your favorite ML frameworks - Pytorch Ignite, Pytorch Lightning, Hugging Face and others. Basic integration guides can be found at [Quick Start](#) section.

In this section we're going to deep-dive into the ways we can extend the basic loggers, manipulate them to track a lot more. The basic loggers can track specific metrics and hyper-params only.

There are two ways Aim callbacks/adapters/loggers can be extended:

- by deriving and overriding the main methods that are responsible for logging.
- by using public property called `experiment` which gives access to underlying `aim.Run` object to easily track new metrics, params and other metadata that would benefit your project.

26.1 Pytorch Ignite

Both callback extension mechanisms are available with Pytorch Ignite. In the example below you'll see how to use the `experiment` property to track confusion matrix as an image using `aim.Image` after the training is completed.

Here is an [example colab notebook](#).

```
from aim import Image
from aim.pytorch_ignite import AimLogger

import matplotlib.pyplot as plt
import seaborn as sns

# Create a logger
aim_logger = AimLogger()
...
@trainer.on(Events.COMPLETED)
def log_confusion_matrix(trainer):
    metrics = val_evaluator.state.metrics
    cm = metrics['cm']
    cm = cm.numpy()
    cm = cm.astype(int)
    classes = ['T-shirt/top', 'Trouser', 'Pullover', 'Dress', 'Coat', 'Sandal', 'Shirt',
    ↪ 'Sneaker', 'Bag', 'Ankle Boot']
    fig, ax = plt.subplots(figsize=(10,10))
    ax = plt.subplot()
    sns.heatmap(cm, annot=True, ax = ax, fmt="d")
    # labels, title and ticks
```

(continues on next page)

(continued from previous page)

```

ax.set_xlabel('Predicted labels')
ax.set_ylabel('True labels')
ax.set_title('Confusion Matrix')
ax.xaxis.set_ticklabels(classes,rotation=90)
ax.yaxis.set_ticklabels(classes,rotation=0)
aim_logger.experiment.track(Image(fig), name='cm_training_end')

```

With Pytorch Ignite there's also a 3rd approach to extend the integration. For example Pytorch Ignite's Tensorboard logger provides a possibility to track model's gradients and weights as histograms. Same can be achieved with Aim

```

from typing import Optional, Union

import torch.nn as nn
from ignite.contrib.handlers.base_logger import BaseWeightsHistHandler
from ignite.engine import Engine, Events

from aim.pytorch_ignite import AimLogger
from aim import Distribution

class AimGradsHistHandler(BaseWeightsHistHandler):
    def __init__(self, model: nn.Module, tag: Optional[str] = None):
        super(GradsHistHandler, self).__init__(model, tag=tag)

    def __call__(self, engine: Engine, logger: AimLogger, event_name: Union[str, Events]) -> None:
        global_step = engine.state.get_event_attrib_value(event_name)
        context = {'subset': self.tag} if self.tag else {}
        for name, p in self.model.named_parameters():
            if p.grad is None:
                continue
            name = name.replace(".", "/")
            logger.experiment.track(
                Distribution(p.grad.detach().cpu().numpy()),
                name=name,
                step=global_step,
                context=context
            )

# Create a logger
aim_logger = AimLogger()

# Attach the logger to the trainer to log model's weights norm after each iteration
aim_logger.attach(
    trainer,
    event_name=Events.ITERATION_COMPLETED,
    log_handler=AimGradsHistHandler(model)
)

```

26.2 Pytorch Lightning

In the [example](#) provided in the Aim GitHub repo using PL + Aim there's already a reference how to customize an integration.

```
def test_step(self, batch, batch_idx):
    ...
    # Track metrics manually
    self.logger.experiment.track(1, name='manually_tracked_metric')
```

So you can track lots of metadata at each iteration of test step: images, texts, whatever is needed by you and supported by Aim.

26.3 Hugging Face

Here is how to extend the basic Hugging Face logger. Below is an example of a CustomCallback that's derived from the AimCallback. The main HF method here is the `on_log()` that's overridden.

This allows us to track any `str` object that is passed to `on_log()` method as `aim.Text`.

```
from aim.hugging_face import AimCallback
from aim import Text

class CustomCallback(AimCallback):
    def on_log(self, args, state, control,
               model=None, logs=None, **kwargs):
        super().on_log(args, state, control, model, logs, **kwargs)

        context = {
            'subset': self._current_shift,
        }
        for log_name, log_value in logs.items():
            if isinstance(log_value, str):
                self.experiment.track(Text(log_value), name=log_name, context=context)
```

26.4 TF/keras

Here is how to track confusion matrices with Aim while extending the default callback provided for `tf.keras`. We have taken and adapted this [example](#). to Aim. Here is how it looks:

```
from aim.tensorflow import AimCallback

class CustomImageTrackingCallback(AimCallback):
    def __init__(self, data):
        super().__init__()
        self.data = data

    def on_epoch_end(self, epoch, logs=None):
        super().on_epoch_end(epoch, logs)
```

(continues on next page)

(continued from previous page)

```

from aim import Image
# Use the model to predict the values from the validation dataset.
test_pred_raw = self.model.predict(test_images)
test_pred = np.argmax(test_pred_raw, axis=1)

# Calculate the confusion matrix.
cm = sklearn.metrics.confusion_matrix(test_labels, test_pred)
# Log the confusion matrix as an image summary.
figure = plot_confusion_matrix(cm, class_names=class_names)
cm_image = Image(figure)

# Log the confusion matrix as an Aim image.
self.experiment.track(cm_image, "Confusion Matrix", step=epoch)

aim_callback = CustomImageTrackingCallback()

model.fit(
    train_images,
    train_labels,
    epochs=5,
    verbose=0, # Suppress chatty output
    callbacks=[aim_callback],
    validation_data=(test_images, test_labels),
)

```

26.5 XGBoost

Here is how to override the AimCallback for XGBoost.

```

from aim import Text
from aim.xgboost import AimCallback

class CustomCallback(AimCallback):

    def after_iteration(self, model, epoch, evals_log):
        for data, metric in evals_log.items():
            for metric_name, log in metric.items():
                self.experiment.track(Text(log), name=metric_name)

        return super().after_iteration(model, epoch, evals_log)

```

26.6 Catboost

Catboost's `.fit()` has `log_cout` parameter which can be used to redirect log output into a custom object which has `write` attribute. Our logger is an object which implements `write` method to parse log string according to its content. Thus, most of the log output will be ignored by our parser logic, but you still can code up your own logic on top of ours to fill the gap for your needs.

```
from aim.catboost import AimLogger

class CustomLogger(AimLogger):

    def write(self, log):
        # Process the log string through our parser
        super().write(log)

        # Do your own parsing
        log = log.strip().split()
        if log[1] == 'bin:':
            value_bin = log[1][4:]
            value_score = self._to_number(log[3])
            self.experiment.track(value_score, name='score')
```

26.7 LightGBM

Here is how to override the `AimCallback` for LightGBM.

```
from aim.lightgbm import AimCallback

class CustomCallback(AimCallback):

    def before_tracking(self, env):
        for item in env.evaluation_result_list:
            # manipulate item here
            pass

    def after_tracking(self, env):
        # do any other action if necessary after tracking value
        pass
```


OVERVIEW

Aim is built around several concepts allowing to make sure that it meets the following criteria:

- **Run data isolation.** Each training run process isolated in terms of data and do not require additional services to run.
- **Scalability.** Aim web app is able to handle 1000s of training runs. Starting from v3.4.0 Aim provides a [Remote Tracking server](#) allowing to run multiple parallel experiments in a distributed multi-host environment.
- **Flexibility.** Aim UI and query language allow users to select, group and filter the tracked data any way they want.

27.1 Aim Components

In order to understand how Aim works, lets take a quick look on a different components it has.

- **Aim Storage.** At its core Aim uses a custom-built storage, based on [rocksdb](#). More details in ‘Where is data collected?’. Data tracked by different training runs collected and indexed in an aim repository (`.aim`). The storage itself is generic; it allows accessing the data as collection of dictionaries and arrays.
- **Aim SDK.** On top of the storage Aim SDK provides functionality to track/select/query data. Additionally, SDK is a layer used by Web APIs and CLI.
- **Aim UI.** Web app allowing to browse run metadata, metrics, images and other tracked data.
- **Aim CLI.** A collection of command line utilities for running Aim web server, managing aim repositories, runs, etc.
- **Remote Tracking server.** A [gRPC](#)-based service accepting incoming traffic and storing data on a centralized server.

The next sections will describe various concepts Aim introduces and provide more detailed look on individual components introduced above.

DATA STORAGE - WHERE AIM DATA IS COLLECTED

This section provides a deep-dive into Aim storage structure. It is important to know the internal storage organization in order to understand how it affects queries performance.

28.1 Storage structure

The core foundation for Aim storage is [rocksdb](#). It is a fast, embedded key-value store maintained by facebook. The aim repository is a collection of individual rocksdb databases with abstraction layers added to manage the collection as one database. The abstraction of a single KV store called Container. Below is the directory structure for a typical Aim project:

```
.aim
  run_metadata.sqlite
  meta/
    index/
    chunks/
      aacf48e769534c32a9cc5a3c/
      80483ab611a24bf5bd8fc288/
      16f83a2c2f50477f8446f322/
    progress/
      aacf48e769534c32a9cc5a3c
      80483ab611a24bf5bd8fc288
    locks/
      aacf48e769534c32a9cc5a3c
      80483ab611a24bf5bd8fc288
  seqs/
    chunks/
      aacf48e769534c32a9cc5a3c/
      80483ab611a24bf5bd8fc288/
      16f83a2c2f50477f8446f322/
    progress/
      aacf48e769534c32a9cc5a3c
      80483ab611a24bf5bd8fc288
    locks/
      aacf48e769534c32a9cc5a3c
      80483ab611a24bf5bd8fc288
```

There are two main parts of the storage:

- `run_metadata.sqlite`: SQLite database for storing Run structured data, such as, creation time, name, Experiment it is attached to, Tags etc.

- **meta/** and **seqs/** directories: a collection of rocksdb storages. Used to write Runs tracked data, such as params, metrics and objects.

In the tree above the hash-strings (i.e. `aacf48e769534c32a9cc5a3c`) represent a single Run. When the new Run is started, aim will create two Containers:

- **Meta** container for logged params as well as metadata about collected sequences, contexts, etc.
- **Sequence** container for the value series.

The reason the actual sequence data separated from metadata is the necessity for fast queries, regardless sequence size.

Additionally, per each container two files will be created to properly manage the container state

- **lock** file, indicating that the container opened in write-mode.
- **progress** file, indicating that container potentially has un-indexed data.

This setup allows implementing concurrent training jobs setup without requiring additional synchronization routines, and without the risk of losing or corrupting the data. For example, two jobs might run on different hosts where aim repo mounted on a shared NFS location.

28.2 What is the index container?

Each run writes data into its own isolated containers. The aim queries require reading the Run metadata from **meta** container. However, with 1000s of runs opening each meta container database will slow-down the queries. Here the indexing of metadata becomes crucial.

Run object maintains lock for both **meta** and **sequence** containers during the training script execution. Run will continue to write its data into its own containers. Once the execution finishes, **meta** container data indexed, the container locks released, and the progress file removed. The **sequence** container data is not indexed, since the individual points of a sequence are not queryable, and the sequence info is available in **meta** container.

28.3 How data written to/read from the storage?

Run object provides interface for logging the Run parameters dict-like data and tracking series of scalars and objects. Aim has a custom encoding layer which translates this hierarchical data into the sets of key-value pairs to be written into rocksdb. The same encoding layer is responsible for re-constructing the tracked data/objects. During query execution, aim SDK will walk through all runs in **index** container + the **meta** containers for chunks which have progress file (remember that progress file indicates potentially un-indexed data for the Run). If the run/sequence match the query expression, the appropriate run/sequence will be yielded. Notice that till this point no data was accessed from the **sequence** container. The sequence data itself is read upon request.

STORAGE INDEXING - HOW AIM DATA IS INDEXED

29.1 Background

When tracking experiment metadata with Aim, each run creates its own isolated space in aim repository. This allows to run multiple concurrent experiments without setting-up additional services responsible for data writes synchronization. Once run is complete, all the data it tracked is being indexed. We call this step run finalization. When the training script terminated with SIGTERM signal, Aim will handle this and make sure that run properly finalized and data is indexed. However, there are cases when training terminated abnormally and data remains unindexed.

29.2 How things worked before?

Due to the chunks of data being unindexed, chunks of data would remain in the runs' separate storage but not in index storage. This means that queries had to open multiple files to read the repo data. Once failed runs started to accumulate, queries will slow down. In order to mitigate this `aim reindex command` has been introduced. The command will scan the aim repo and index all stalled runs.

29.3 Automatic indexing

Though `aim reindex` command will address the performance issues it is not the most convenient way to do. The questions such as “When should I run `aim reindex`?” or “How frequent should I run `aim reindex`?” depend on the actual aim repository and use-case. Thus, we need to automate the indexing of aim repository. Each time `aim up` command is ran, Aim will spawn a background thread along with the web server. The thread will check for the unindexed runs and reindex them one at the time. This will keep queries performance high without locking the index storage for too long.

29.4 Conclusion

With the new automatic indexing logic in place, users don't have to manually run `aim reindex` command. It is still in place for cases when all the runs data should be indexed at once. The combination of automatic (implicit) and manual (explicit) reindexing makes sure aim repo has good performance in a long-term usage scenarios and provides good overall user experience.

CONCEPTS

All the functionality in Aim is build around several key concepts. This chapter will give a brief overview of these core concepts. For more details please check the [Reference](#) section or [Glossary](#).

30.1 Aim Run

Run is an abstraction representing the tracked data for a single experiment. Its in memory model is SDK class `aim.Run`. It is a core class used in your training script for tracing metrics and objects, as well as storing training hyperparams and other data. Run object are queryable and UI provides a rich functionality for exploring runs and browsing single run details.

30.2 Aim Repo

While you do multiple training experiments, multiple runs data stored in a single directory called Aim repository (repo for short). You can think of aim repo as an application centralized database. SDK provides an in memory model for repo `aim.Repo` class. It is responsible for repository resources management and might be used to query and/or iterate over the stored data.

30.3 Run Params

Each run has a set of parameters associating with it. This might include the training script hyperparameters, dataset information, etc. The Run object provides dictionary-like interface to set and access run params. Run parameters are also available in the context of queries. You can set the whole configuration at once with the syntax like this:

```
run['hparams'] = conf
```

At this moment Run supports setting configuration from Python dictionaries and `OmegaConf` configs. Support of popular configuration formats constantly added. You can check the full list in [Supported Data types](#) section.

30.4 Run Sequence

The sequence is a set of homogeneous ordered objects. In aim sequence must be bound to the Run object. When the value is tracked in aim, it is appended to an existing or newly created Sequence object. The entire sequences can be queried using aim QL and each sequence can be sliced further down. Sequence object is agnostic to the element type it holds. The way how the sequence represented in UI, and the set of additional operations it might have depends on the element type. For example Metric is a sequence of scalars. It can be represented as a value chart in UI, and SDK provides methods to convert it to `numpy.ndarray`.

30.5 Sequence Context

The sequence context provides a mechanism to query/group multiple sequences beyond simple string comparison on sequence name. Sequences with the same name but with different context can perfectly coexist in the scope of one Run. In other words, sequence defined by its Run, name and context. The example usage of this is tracking the same metric 'loss' for different stages of training (train, validation, test). The resulting Run will have 3 metrics: 1) 'loss' {'subset': 'train'} 2) 'loss' {'subset': 'val'} 3) 'loss' {'subset': 'test'} Here is a small code example demonstrating how to specify context for a metric sequence:

```
from aim import Run

aim_run = Run()
for i in range(100):
    if i % 2 == 0:
        aim_run.track(i, name=r'numbers', context={'odds': True})
    else:
        aim_run.track(i, name=r'numbers', context={'odds': False})
```


RUNNING AIM WITH PROFILING

Aim comes with profiling feature which is logging all api requests to the backend into a directory inside your repository. We use `pyinstrument` as underlying profiler. To toggle profiling, run `aim up` with `--profiler` flag.

```
$ aim up --profiler
```

This will instruct backend to create a new directory inside your repository (`.aim/profiler`). On every api call, profiler will create an `.html` file containing whole run trace of that api. Basically you can navigate into that directory and open the file in your browser to see where's the performance bottleneck.

31.1 Why would you need to enable the profiling.

Well you don't, unless we explicitly ask you to in case you have any performance issues while using Aim.

Please note that no data is sent to us when profiling is toggled. Everything is stored locally and managed by the end user.

STORING METRIC SEQUENCE DATA FOR FAST ACCESS

32.1 Background

Aim allows tracking metric sequences without any limitation of the number of steps tracked. With a few thousands of steps per metric it is OK to load the entire sequence into memory and slice it to get the desired number of points. However, when many (tens of millions) steps are tracked, this approach cannot be applied. Starting from version 3.11 Aim metric sequence data format has changed to improve the data collection/sampling performance for long sequences.

32.2 What has changed?

Each value tracked has its unique key in the Run's rocksdb database. Previously the tracking step has been used as the part of this key. This means that in order to get the sample of K values we have to either do K random access operations or iterate over entire sequence and skip values. Both have negative implication on data read performance. With Aim 3.11 the key uses the stable random hashing algorithm instead of tracking step. Since the keys are sorted, reading first K keys is a good approximation for the entire metric sequence.

32.3 What to do if the Metric has been logged already?

In order to speed-up the metrics read for existing aim Repos the data format upgrade should be executed. In cases when Run has been opened for write, the upgrade is done automatically, since it is not possible to mix two formats. For the rest of the cases new CLI command has been added to apply data format change for selection of Runs or for the entire repo:

```
# update selected list of runs
aim storage upgrade 3.11+ <RUN_HASH_1> <RUN_HASH_2> ...

# update all runs in the repo
aim storage --repo <REPO_PATH> upgrade 3.11+ '*'
```

The Run data backup with old format is available in case of any issues. In order to rollback format change and get back to the old format run the following command:

```
aim storage restore <RUN_HASH>
```


TRACK AND COMPARE GANS WITH AIM

33.1 Overview

[Generative Adversarial Networks](#), or GANs, are deep-learning-based generative models.

Generative modeling is an unsupervised learning task in machine learning that involves automatically discovering and learning the patterns of input data in such a way that the model can be used to generate new examples that plausibly could have been drawn from the original dataset.

In this guide we will show you how to integrate Aim with your GAN and GAN with EMA to compare the generated images from both experiments to compare their performances.

33.2 Experiment

We will train and compare a regular GAN vs GAN with EMA. EMA is a technique for parameter averaging in GAN training, which computes an exponentially discounted sum of weights.

We will use [lightweight-gan](#) model implemented by [lucidrains](#) and [MetFaces Dataset](#) as a training dataset.

To be able to analyze the results we will fix random 64 points and track them during the training both for a regular GAN and GAN w/ EMA.

33.3 Track images with Aim

1. Initialize a new run in the trainer class to collect and store sequences of images:

```
class Trainer():
    def __init__(
        self,
        name = 'default',
        results_dir = 'results',
        models_dir = 'models',
        ...
    ):
        ...
    self.run = aim.Run()           # Initialize aim.Run
    self.run['hparams'] = hparams  # Log hyperparams
    ...
```

Code on [GitHub](#)

1. Track images generated by a regular GAN:

```
# Regular GAN

# Get generated images
generated_images = self.generate_(self.GAN.G, latents)

aim_images = []
for idx, image in enumerate(generated_images):
    ndarr = image.mul(255).add_(0.5).clamp_(0, 255).permute(1, 2, 0).to('cpu', torch.
↳ uint8).numpy()
    im = PIL.Image.fromarray(ndarr)
    aim_images.append(aim.Image(im, caption=f'#{idx}'))

# Store with Aim (name="generated" and context.ema=0)
self.run.track(value=aim_images, name='generated', step=self.steps, context={'ema':
↳ False})
```

Code on [GitHub](#)

1. Track images generated by a GAN with enabled EMA:

```
# GAN with moving averages

# Get generated images
generated_images = self.generate_(self.GAN.GE, latents)

aim_images = []
for idx, image in enumerate(generated_images):
    ndarr = image.mul(255).add_(0.5).clamp_(0, 255).permute(1, 2, 0).to('cpu', torch.
↳ uint8).numpy()
    im = PIL.Image.fromarray(ndarr)
    aim_images.append(aim.Image(im, caption=f'EMA #{idx}'))

# Store with Aim (name="generated" and context.ema=1)
self.run.track(value=aim_images, name='generated', step=self.steps, context={'ema': True}
↳ )
```

Code on [GitHub](#)

33.4 Explore the results with Aim UI

1. Visualize images generated by a regular GAN:

1. Visualize images generated by GAN with EMA:

As you may notice GAN with EMA converges in an exponential fashion and has better results at the end.

1. Let's compare the final step of the two methods side by side:

33.5 Conclusion

As you can see GAN with EMA performed much better compared to the regular one.

With Aim you can easily compare diff groups of tracked images from diff runs.

Group them by the run hash, other parameters available to slice and dice and observe the difference between the runs.

AIM CLI

Aim CLI offers a simple interface to easily organize and record your experiments. Paired with the Python Library, Aim is a powerful utility to record, search and compare AI experiments. Here are the set of commands supported:

Command	Description
<code>init</code>	Initialize the aim repository.
<code>version</code>	Displays the version of aim cli currently installed.
<code>up</code>	Runs Aim web UI for the given repo.
<code>reindex</code>	Process runs left in 'in progress' state and optimized finished runs.
<code>server</code>	Run aim remote tracking server accepting incoming RPC requests.
<code>runs</code>	Manage run data for the given repo.
<code>convert</code>	Tool-set for converting 3rd party data into Aim readable format.
<code>storage</code>	Maintain/update Aim repository internal data formats.

34.1 init

****This step is optional.**** Initialize the aim repo to record the experiments.

```
$ aim init
```

Creates `.aim` directory to save the recorded experiments to. Running `aim init` in an existing repository will prompt the user for re-initialization.

Args	Description
<code>--repo <repo_path></code>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>

****Beware:**** Re-initialization of the repo clears `.aim` folder from previously saved data and initializes new repo.

****Note:**** This command is not necessary to be able to get started with Aim as aim is automatically initializes with the first aim function call.

34.2 version

Display the Aim version installed.

```
$ aim version
```

34.3 up

Start the Aim web UI locally.

```
$ aim up [ARGS]
```

Args	Description
<code>-h</code> | <code>--host</code> <host>	Specify host address.
<code>-p</code> | <code>--port</code> <port>	Specify port to listen to.
<code>--repo</code> <repo_path>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>
<code>--dev</code>	Run UI in development mode.
<code>--profiler</code>	Enables API profiling which logs run trace inside <code>.aim/profiler</code> directory.
<code>--log-level</code>	Specifies log level for python logging package. <code>WARNING</code> by default, <code>DEBUG</code> when <code>--dev</code> option is provided.

34.4 server

Run a gRPC server to collect tracked data from remote clients.

```
$ aim server [ARGS]
```

Args	Description
<code>--repo</code> <repo_path>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>
<code>-h</code> | <code>--host</code> <host>	Specify host address.
<code>-p</code> | <code>--port</code> <port>	Specify port to listen to. <i>Default is 53800.</i>
<code>-w</code> | <code>--workers</code> <N>	Specify number of gPRC workers. <i>Default is 1 worker.</i>
<code>--ssl-keyfile</code>	Specify path to keyfile for secure connection.
<code>--ssl-certfile</code>	Specify path to cert. file for secure connection.
<code>--log-level</code>	Specifies log level for python logging package. <code>WARNING</code> by default.

34.5 runs

Upgrade Aim repository runs data.

```
$ aim runs [ARGS] SUBCOMMAND
```

Args	Description
<code>--repo <repo_path></code>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>

runs subcommands

Sub-command	Description
<code>ls</code>	List runs in aim repository.
<code>rm</code>	Remove run data for given run hashes. At least one run should be specified.
<code>cp</code>	Copy run data for given run hashes. At least one run should be specified.
<code>mv</code>	Move run data for given run hashes. At least one run should be specified.
<code>upload</code>	Create a snapshot of <code>.aim</code> directory in cloud. Bucket name should be specified.

Global expression (*) support is available for run hashes. If hash contains *, it must be enclosed within quotes (') as bash resolves the expression before passing it to `aim runs` command.

```
$ aim runs ls
```

```
$ aim runs rm [HASH] ...
```

```
$ aim runs cp [ARGS] [HASH] ...
```

Args	Description
<code>--destination <dest_repo_path></code>	Path to destination repo. Required.

```
$ aim runs mv [ARGS] [HASH] ...
```

Args	Description
<code>--destination <dest_repo_path></code>	Path to destination repo. Required.

```
$ aim runs upload [ARGS] ...
```

Args	Description
<code>--bucket <bucket_name></code>	Name of the cloud storage bucket. Required.

34.6 convert

Tool-set for converting 3rd party data into Aim readable format.

```
$ aim convert [ARGS] SUBCOMMAND
```

Args	Description
<code>--repo <repo_path></code>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>

convert subcommands

Sub-command	Description
<code>tensorboard</code>	Convert from Tensorboard logs.
<code>mlflow</code>	Convert from MLFlow logs.

Sub-command: tensorboard

Options	Description
<code>--flat</code>	Disregard context directory and treat them as distinct run directories. Inactive by default.
<code>--no-cache</code>	Ignore previously cached results and process the logs entirely. Disabled by default.

Sub-command: mlflow

Options	Description
<code>--tracking_uri</code>	MLFlow logs URI. Can be either an HTTP/HTTPS URI for a remote server, a database connection string, or a local path.
<code>-e</code> <code>&#124;</code> <code>--experiment</code>	MLFlow Experiment name. If specified, only runs for <code>exp_name</code> will be converted.

34.7 storage

Perform various maintenance operations on Aim repository.

```
$ aim storage [ARGS] SUBCOMMAND
```

Args	Description
<code>--repo <repo_path></code>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>

storage subcommands

Sub-command	Description
<code>upgrade 2to3</code>	Upgrades legacy Aim repository from 2.x to 3.0.
<code>upgrade 3.11+</code>	Update metric sequence data format for given runs. At least one run should be specified.
<code>restore</code>	Rollback Run to old metric format if run backup is available.
<code>reindex</code>	Update index to include all runs in Aim repo which are left in progress.
<code>prune</code>	Remove dangling params/sequences with no referring runs.

Sub-command: update 2to3

```
$ aim storage upgrade 2to3 [ARGS]
```

Args	Description
--skip-failed-runs	Use this flag to skip runs which are failed/have missing or incomplete data.
--skip-checks	Use this flag to skip new repository consistency checks.
--drop-existing	Use this flag to clear old .aim directory. By default old data is kept in .aim_legacy.

Sub-command: update 3.11+

```
$ aim storage upgrade 3.11+ [HASH] ...
```

Sub-command: restore

```
$ aim storage restore [HASH] ...
```

Sub-command: reindex | Args | Description | |
 | --finalize-only | Only finalize runs left in ‘in progress’ state. Do not attempt runs optimization. |

Sub-command: prune

```
$ aim storage prune
```


AIM STATUS WATCHER CLI

Aim status CLI offers an interface to start training runs status monitoring and configure how the notifications should be received. The entry point to this CLI is `aim-watcher`:

```
$ aim-watcher [ARGS] SUBCOMMAND
```

Args	Description
<code>--repo <repo_path></code>	Path to parent directory of <code>.aim</code> repo. <i>Current working directory by default.</i>

Sub-Commands	Description
<code>start</code>	Start watcher service to monitor and report stuck/failed Runs.
<code>notifiers</code>	Configure how notifications should be received.

35.1 start

Start watcher service for the given aim Repo.

```
aim-watcher --repo . start
```

35.2 notifiers

Manipulate notifiers configuration.

notifiers subcommands

```
$ aim-watcher notifiers add
```

```
$ aim-watcher notifiers list
```

```
$ aim-watcher notifiers remove [NOTIFIER_ID]
```

```
$ aim-watcher notifiers disable [NOTIFIER_ID]
```

```
$ aim-watcher notifiers enable [NOTIFIER_ID]
```


36.1 aim.sdk.repo module

class `aim.sdk.repo.Repo`(*path*, *, *read_only=None*, *init=False*)

Aim repository object.

Provides methods for repositories creation/opening/cleanup. Provides APIs for accessing Runs. Provides API for querying Runs/Metrics based on a given expression.

Parameters

- **path** (`str`) – Path to Aim repository.
- **read_only** (`bool`, optional) – Flag for opening Repo in readonly mode. False by default.
- **init** (`bool`, optional) – Flag used to initialize new Repo. False by default. Recommended to use `aim init` command instead.

collect_params_info()

Utility function for getting run meta-parameters.

Returns

All runs meta-parameters.

Return type

`dict`

collect_sequence_info(*sequence_types*)

Utility function for getting sequence names and contexts for all runs by given sequence types.

Parameters

sequence_types (`tuple[str]`, optional) – Sequence types to get tracked sequence names/contexts for. Defaults to 'metric'.

Returns

Tree of sequences and their contexts grouped by sequence type.

Return type

`dict`

copy_runs(*run_hashes*, *dest_repo*)

Copy multiple Runs data from current aim repository to destination aim repository

Parameters

- **run_hashes** (`str`) – list of Runs to be copied.
- **dest_repo** ([Repo](#)) – destination Repo instance to copy Runs

Returns

(True, []) if all runs were copied successfully, (False, list) with list of remaining runs otherwise.

classmethod **default_repo**(*init=False*)

Named constructor for default repository.

Searches nearest *.aim* directory from current directory to roo directory. If not found, return Repo for current directory.

Parameters

init (bool, optional) – Flag used to initialize new Repo. False by default. Recommended to use *aim init* command instead.

Returns

Repo object.

delete_run(*run_hash*)

Delete Run data from aim repository

This action removes run data permanently and cannot be reverted. If you want to archive run but keep it's data use *repo.get_run(run_hash).archived = True*.

Parameters

run_hash (str) – Run to be deleted.

Returns

True if run deleted successfully, False otherwise.

delete_runs(*run_hashes*)

Delete multiple Runs data from aim repository

This action removes runs data permanently and cannot be reverted. If you want to archive run but keep it's data use *repo.get_run(run_hash).archived = True*.

Parameters

run_hashes (str) – list of Runs to be deleted.

Returns

(True, []) if all runs deleted successfully, (False, list) with list of remaining runs otherwise.

classmethod **exists**(*path*)

Check Aim repository existence.

Parameters

path (str) – Path to Aim repository.

Returns

True if repository exists, False otherwise.

classmethod **from_path**(*path, read_only=None, init=False*)

Named constructor for Repo for given path.

Parameters

- **path** (str) – Path to Aim repository.
- **read_only** (bool, optional) – Flag for opening Repo in readonly mode. False by default.
- **init** (bool, optional) – Flag used to initialize new Repo. False by default. Recommended to use *aim init* command instead.

Returns*Repo* object.**get_run**(*run_hash*)

Get run if exists.

Parameters**run_hash** (str) – Run hash.**Returns**Run object if hash is found in repository. *None* otherwise.**iter_runs**()

Iterate over Repo runs.

Yields

next Run in readonly mode .

move_runs(*run_hashes*, *dest_repo*)

Move multiple Runs data from current aim repository to destination aim repository

Parameters

- **run_hashes** (str) – list of Runs to be moved.
- **dest_repo** (*Repo*) – destination Repo instance to move Runs

Returns

(True, []) if all runs were moved successfully, (False, list) with list of remaining runs otherwise.

query_audios(*query*="", *report_mode*=*QueryReportMode.PROGRESS_BAR*)

Get audio collections satisfying query expression.

Parameters

- **query** (str) – query expression.
- **report_mode** (*QueryReportMode*, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). *QueryReportMode.PROGRESS_BAR* if not specified.

Returns

Iterable for audio sequences matching query expression.

Return type*SequenceCollection***query_distributions**(*query*="", *report_mode*=*QueryReportMode.PROGRESS_BAR*)

Get distribution collections satisfying query expression.

Parameters

- **query** (str) – query expression.
- **report_mode** (*QueryReportMode*, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). *QueryReportMode.PROGRESS_BAR* if not specified.

Returns

Iterable for distribution sequences matching query expression.

Return type*SequenceCollection*

query_figure_objects(*query*="", *report_mode*=*QueryReportMode.PROGRESS_BAR*)

Get Figures collections satisfying query expression.

Parameters

- **query** (str) – query expression.
- **report_mode** (*QueryReportMode*, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). *QueryReportMode.PROGRESS_BAR* if not specified.

Returns

Iterable for Figure sequences matching query expression.

Return type

SequenceCollection

query_images(*query*="", *report_mode*=*QueryReportMode.PROGRESS_BAR*)

Get image collections satisfying query expression.

Parameters

- **query** (str) – query expression.
- **report_mode** (*QueryReportMode*, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). *QueryReportMode.PROGRESS_BAR* if not specified.

Returns

Iterable for image sequences matching query expression.

Return type

SequenceCollection

query_metrics(*query*="", *report_mode*=*QueryReportMode.PROGRESS_BAR*)

Get metrics satisfying query expression.

Parameters

- **query** (str) – query expression.
- **report_mode** (*QueryReportMode*, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). *QueryReportMode.PROGRESS_BAR* if not specified.

Returns

Iterable for metrics matching query expression.

Return type

SequenceCollection

query_runs(*query*="", *paginated*=*False*, *offset*=*None*, *report_mode*=*QueryReportMode.PROGRESS_BAR*)

Get runs satisfying query expression.

Parameters

- **query** (str, optional) – query expression. If not specified, query results will include all runs.
- **paginated** (bool, optional) – query results pagination flag. False if not specified.
- **offset** (str, optional) – *hash* of Run to skip to.

- **report_mode** (QueryReportMode, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). QueryReportMode.PROGRESS_BAR if not specified.

Returns

Iterable for runs/metrics matching query expression.

Return type

SequenceCollection

query_texts(*query*="", *report_mode*=QueryReportMode.PROGRESS_BAR)

Get text collections satisfying query expression.

Parameters

- **query** (str) – query expression.
- **report_mode** (QueryReportMode, optional) – indicates report mode (0: DISABLED, 1: PROGRESS BAR, 2: PROGRESS TUPLE). QueryReportMode.PROGRESS_BAR if not specified.

Returns

Iterable for text sequences matching query expression.

Return type

SequenceCollection

classmethod rm(*path*)

Remove Aim repository.

Parameters

path (str) – Path to Aim repository.

36.2 aim.sdk.run module

class aim.sdk.run.Run(*args, **kwargs)

Run object used for tracking metrics.

Provides method [track](#) to track value and object series for multiple names and contexts. Provides dictionary-like interface for Run object meta-parameters. Provides API for iterating through tracked sequences.

Parameters

- **run_hash** (str, optional) – Run’s hash. If skipped, generated automatically.
- **repo** (Union[Repo, str], optional) – Aim repository path or Repo object to which Run object is bound. If skipped, default Repo is used.
- **read_only** (bool, optional) – Run creation mode. Default is False, meaning Run object can be used to track metrics.
- **experiment** (str, optional) – Sets Run’s *experiment* property. ‘default’ if not specified. Can be used later to query runs/sequences.
- **force_resume** (bool, optional) – Forcefully resume stalled Run.
- **system_tracking_interval** (int, optional) – Sets the tracking interval in seconds for system usage metrics (CPU, Memory, etc.). Set to *None* to disable system metrics tracking.
- **log_system_params** (bool, optional) – Enable/Disable logging of system params such as installed packages, git info, environment variables, etc.

__delitem__(*key*)

Remove key from run meta-params. :param key: meta-parameter path

__getitem__(*key*)

Get run meta-parameter by key.

Parameters**key** – path to Run meta-parameter.**Returns**

Collected sub-tree of Run meta-parameters.

Examples

```
>>> run = Run('3df703c')
>>> run['hparams'] # -> {'batch_size': 42}
>>> run['hparams', 'batch_size'] # -> 42
```

__setitem__(*key, val*)

Set Run top-level meta-parameter.

Parameters

- **key** (str) – Top-level meta-parameter name. Use ellipsis to reset run's all meta-parameters.
- **val** – Meta-parameter value.

Examples

```
>>> run = Run('3df703c')
>>> run[...] = params
>>> run['hparams'] = {'batch_size': 42}
```

add_tag(*value*)

Add tag to run

Parameters**value** (str) – Tag to add.**collect_sequence_info**(*sequence_types, skip_last_value=False*)

Retrieve Run's all sequences general overview.

Parameters

- **sequence_types** – Type names of sequences for which to collect name/context pairs.
- **skip_last_value** (bool, optional) – Boolean flag to include tracked sequence last value in sequence info. False by default.

Returnslist of sequence's *context, name* and optionally last tracked value triplets.**Return type**

list

dataframe(*include_props=True, include_params=True*)

Get run properties and params as pandas DataFrame

Parameters

- **include_props** – (int, optional): If true, include run structured props
- **include_params** – (int, optional): If true, include run parameters

get_audio_sequence(*name, context*)

Retrieve audios sequence by its name and context.

Parameters

- **name** (str) – Tracked audios sequence name.
- **context** (Context) – Tracking context.

Returns

Audios object if exists, *None* otherwise.

get_distribution_sequence(*name, context*)

Retrieve distributions sequence by it's name and context.

Parameters

- **name** (str) – Tracked distribution sequence name.
- **context** (Context) – Tracking context.

Returns

Distributions object if exists, *None* otherwise.

get_figure_sequence(*name, context*)

Retrieve figure sequence by its name and context.

Parameters

- **name** (str) – Tracked figure sequence name.
- **context** (Context) – Tracking context.

Returns

Figures object if exists, *None* otherwise.

get_image_sequence(*name, context*)

Retrieve images sequence by it's name and context.

Parameters

- **name** (str) – Tracked image sequence name.
- **context** (Context) – Tracking context.

Returns

Images object if exists, *None* otherwise.

get_log_records()

Retrieve duplicated terminal logs for a run

Returns

Sequence object if exists, *None* otherwise.

get_metric(*name*, *context*)

Retrieve metric sequence by it's name and context.

Parameters

- **name** (str) – Tracked metric name.
- **context** (Context) – Tracking context.

Returns

Metric object if exists, *None* otherwise.

get_terminal_logs()

Retrieve duplicated terminal logs for a run

Returns

Logs object if exists, *None* otherwise.

get_text_sequence(*name*, *context*)

Retrieve texts sequence by it's name and context.

Parameters

- **name** (str) – Tracked text sequence name.
- **context** (Context) – Tracking context.

Returns

Texts object if exists, *None* otherwise.

iter_metrics_info()

Iterator for all run metrics info.

Yields

tuples of (name, context, run) where run is the Run object itself and name, context defines Metric type sequence (with values of *float* and *int*).

iter_sequence_info_by_type(*dtypes*)

Iterator for run sequence infos for the given object data types

Parameters

dtypes – The objects data types list.

Yields

tuples of (name, context, run) where run is the Run object itself and name, context defines sequence for one of *dtypes* types.

metrics()

Get iterable object for all run tracked metrics.

Returns

Iterable for run metrics.

Return type

MetricCollection

Examples

```
>>> run = Run('3df703c')
>>> for metric in run.metrics():
>>>     metric.values.sparse_numpy()
```

remove_tag(tag_name)

Remove run tag.

Parameters

tag_name (str) – name of tag to be removed.

report_progress(*, expect_next_in=0, block=False)

Report progress for the run. Report the expected time for the next progress report.

If no progress reports are received by the expiry date (plus the grace period), the run is considered as failed.

Parameters

- **expect_next_in** – (int, optional): The number of seconds to wait before the next progress report.
- **block** – (bool, optional): If true, block the thread until the report is written to filesystem.

report_successful_finish(*, block=True)

Report successful finish of the run. If the run is not marked as successfully finished, it can potentially be considered as failed.

Parameters

block – (bool, optional): If true, block the thread until the report is written to filesystem.

track(value, name=None, step=None, epoch=None, *, context=None)

Main method for tracking numeric value series and object series.

Parameters

- **value** – The tracked value.
- **name** (str) – Tracked sequence name.
- **step** (int, optional) – Sequence tracking iteration. Auto-incremented if not specified.
- **epoch** (int, optional) – The training epoch.
- **context** (dict, optional) – Sequence tracking context.

Appends the tracked value to sequence specified by *name* and *context*. Appended values should be of the same type, in other words, sequence is a homogeneous collection.

property active

Check if run is active or not.

Getter

Returns run's active state.

Type

bool

property archived

Check is run archived or not.

Getter

Returns run's archived state.

Setter

Archive/un-archive run.

Type

bool

property creation_time

Run object creation time [UTC] as timestamp.

Getter

Returns run creation time.

property description

Run description, set by user.

Getter

Returns run's description.

Setter

Sets run's description.

Type

string

property duration

Run duration in seconds (end_time - creation_time)

Getter

Returns run duration.

property end_time

Run finalization time [UTC] as timestamp.

Getter

Returns run finalization time.

property experiment

Run experiment.

Getter

Returns run's experiment name.

Setter

Sets run's experiment.

Type

string

property name

Run name, set by user.

Getter

Returns run's name.

Setter

Sets run's name.

Type

string

property tags

List of run tags.

Getter

Returns run's tag list.

36.3 aim.sdk.objects.image

class aim.sdk.objects.image.**Image**(*args, **kwargs)

Image object used to store image objects in Aim repository...

Parameters

- **((image) – obj):** pillow *Image* object or *torch.Tensor* or *numpy.array* used to construct *aim.Image*.
- **caption** (str, optional) – Optional image caption. “” by default.
- **format** (*For more information on the*) – Parameter for PIL's .save() method. 'png' by default.
- **quality** (int, optional) – Parameter for PIL's .save() method. 90 by default.
- **optimize** (bool, optional) – Parameter for PIL's .save() method. False by default.
- **format** –
- **parameters** (*quality and optimize*) –
- **documentation.** (*refer to PIL*) –

Example of params to reduce quality of the image:

format='jpeg', optimize=True, quality=85

json()

Dump image metadata to a dict

to_pil_image()

Method to convert aim.Image to pillow Image

property caption

Image caption, set by user.

Getter

Returns image caption.

Setter

Sets image caption.

Type

string

property format

Stored image format.

Getter

Returns image format.

Type

string

property height

Stored image height.

Getter

Returns image height.

Type

string

property size

Stored image size.

Getter

Returns image (width, height) pair.

Type

string

property width

Stored image width.

Getter

Returns image width.

Type

string

36.4 aim.sdk.objects.distribution

class aim.sdk.objects.distribution.Distribution(*args, **kwargs)

Distribution object used to store distribution objects in Aim repository.

Parameters

- ((*hist*) – obj:): Array-like object of data sampled from a distribution.
- **bin_count** (int, optional) – Optional distribution bin count for binning *samples*. 64 by default, max 512.
- (– obj:, optional): Array-like object representing bin frequency counts. Max 512 bins allowed. *samples* must not be specified.
- **bin_range** (tuple, optional) – Tuple of (start, end) bin range.

classmethod from_histogram(*hist*, *bin_range*)

Create Distribution object from histogram.

Parameters

- ((*hist*) – obj:): Array-like object representing bin frequency counts. Max 512 bins allowed.
- **bin_range** (tuple, optional) – Tuple of (start, end) bin range.

classmethod from_samples(*samples*, *bin_count*=64)

Create Distribution object from data samples.

Parameters

- ((*samples*) – obj:): Array-like object of data sampled from a distribution.

- **bin_count** (int, optional) – Optional distribution bin count for binning *samples*. 64 by default, max 512.

json()

Dump distribution metadata to a dict

to_np_histogram()

Return *np.histogram* compatible format of the distribution

property bin_count

Stored distribution bin count

Getter

Returns distribution bin_count.

Type

string

property range

Stored distribution range

Getter

Returns distribution range.

Type

List

property ranges

Stored distribution ranges

Getter

Returns distribution ranges as *np.array*.

Type

np.ndarray

property weights

Stored distribution weights

Getter

Returns distribution weights as *np.array*.

Type

np.ndarray

36.5 aim.sdk.objects.audio

class aim.sdk.objects.audio.**Audio**(*args, **kwargs)

Audio object used to store audio objects in Aim repository..

Currently, audio formats are limited to mp3, wav, flac

Parameters

- **data** – file path, bytes, io.BaseIO or numpy.array (only for WAV)
- **format** (str) – Format of the audio source
- **rate** (int) – Rate of the audio file, for WAV defaults to 22500

- **caption** (str, optional) – Optional audio caption. “” by default.

get()

Reads data from the inner container and writes it to a buffer

Returns: io.BytesIO

json()

Dump audio metadata to a dict

to_numpy()

This method converts WAV to Numpy array. Other audio formats are not supported at this moment.

Returns: numpy array

36.6 aim.sdk.objects.text

class aim.sdk.objects.text.**Text**(*args, **kwargs)

Text object used to store text objects in Aim repository.

Parameters

(*text*) – obj: str object used to construct *aim.Text*.

36.7 aim.sdk.objects.figure

class aim.sdk.objects.figure.**Figure**(*args, **kwargs)

Figure object can be used for storing Plotly or Matplotlib figures into Aim repository. Core functionality is based on Plotly.

Parameters

(*obj*) – obj: plotly or matplotlib figure object.

json()

Dump figure metadata to a dict

36.8 aim.sdk.sequence module

class aim.sdk.sequence.**Sequence**(name, context, run)

Class representing single series of tracked value.

Objects series can be retrieved as Sequence regardless the object’s type, but subclasses of Sequence might provide additional functionality. Provides interface to access tracked values, steps, timestamps and epochs. Values, epochs and timestamps are accessed via [aim.storage.arrayview.ArrayView](#) interface.

classmethod **allowed_dtypes()**

classmethod to get allowed object types for particular sequence

For example, numeric sequences a.k.a. Metric allow float and integer numbers. The base Sequence allows any value, and to indicate that, *allowed_dtypes* returns “*”.

first_step()

Get sequence tracked first step.

Required to implement ranged and sliced data fetching.

last_step()

Get sequence tracked last step.

Required to implement ranged and sliced data fetching.

classmethod sequence_name()

classmethod to get retrieve sequence's registered name

property epochs

Tracked epochs array as `ArrayView`.

Getter

Returns epochs `ArrayView`.

property timestamps

Tracked timestamps array as `ArrayView`.

Getter

Returns timestamps `ArrayView`.

property values

Tracked values array as `ArrayView`.

Getter

Returns values `ArrayView`.

36.9 aim.sdk.sequences.metric module

36.10 aim.sdk.sequences.image_sequence module

class `aim.sdk.sequences.image_sequence.Images(name, context, run)`

Class representing series of Image objects or Image lists.

classmethod sequence_name()

classmethod to get retrieve sequence's registered name

36.11 aim.sdk.sequences.distribution_sequence module

class `aim.sdk.sequences.distribution_sequence.Distributions(name, context, run)`

Class representing series of Distribution objects.

classmethod sequence_name()

classmethod to get retrieve sequence's registered name

36.12 aim.sdk.sequences.audio_sequence module

class aim.sdk.sequences.audio_sequence.**Audios**(*name, context, run*)

Class representing series of Audio objects or Audio lists.

classmethod **sequence_name**()

classmethod to get retrieve sequence's registered name

36.13 aim.sdk.sequences.text_sequence module

class aim.sdk.sequences.text_sequence.**Texts**(*name, context, run*)

Class representing series of Text objects.

classmethod **sequence_name**()

classmethod to get retrieve sequence's registered name

36.14 aim.sdk.sequences.figure_sequence module

class aim.sdk.sequences.figure_sequence.**Figures**(*name, context, run*)

Class representing series of Plotly figure objects or Plotly lists.

classmethod **sequence_name**()

classmethod to get retrieve sequence's registered name

36.15 aim.sdk.sequence_collection module

class aim.sdk.sequence_collection.**SequenceCollection**

Abstract interface for collection of tracked series/sequences.

Typically represents sequences of a same run or sequences matching given query expression.

abstract **iter**()

Get Sequence iterator for collection's sequences.

Yields

Next sequence object based on implementation.

abstract **iter_runs**()

Get SequenceCollection iterator for collection's runs.

Yields

Next run's SequenceCollection based on implementation.

class aim.sdk.sequence_collection.**SingleRunSequenceCollection**(*run, seq_cls=<class 'aim.sdk.sequence.Sequence'>, query="", runs_proxy_cache=None, timezone_offset=0*)

Bases: [*SequenceCollection*](#)

Implementation of SequenceCollection interface for a single Run.

Method *iter()* returns Sequence iterator which yields Sequence matching query from run's sequences. Method *iter_runs()* raises StopIteration, since the collection is bound to a single Run.

Parameters

- **run** (Run) – Run object for which sequences are queried.
- **seq_cls** (type) – The collection's sequence class. Sequences not matching to `seq_cls.allowed_dtypes` will be skipped. *Sequence* by default, meaning all sequences will match.
- **query** (str, optional) – Query expression. If specified, method *iter()* will return iterator for sequences matching the query. If not, method *iter()* will return iterator for run's all sequences.

```
class aim.sdk.sequence_collection.QuerySequenceCollection(repo, seq_cls=<class
                                                         'aim.sdk.sequence.Sequence'>, query="",
                                                         re-
                                                         port_mode=QueryReportMode.PROGRESS_BAR,
                                                         timezone_offset=0)
```

Bases: [SequenceCollection](#)

Implementation of SequenceCollection interface for repository's sequences matching given query.

Method *iter()* returns Sequence iterator, which yields Sequence matching query from currently iterated run's sequences. Once there are no sequences left in current run, repository's next run is considered. Method *iter_runs()* returns SequenceCollection iterator for repository's runs.

Parameters

- **repo** (Repo) – Aim repository object.
- **seq_cls** (type) – The collection's sequence class. Sequences not matching to `seq_cls.allowed_dtypes` will be skipped. *Sequence* by default, meaning all sequences will match.
- **query** (str, optional) – Query expression. If specified, method *iter()* will skip sequences not matching the query. If not, method *iter()* will return iterator for all sequences in repository (that's a lot of sequences!).

```
class aim.sdk.sequence_collection.QueryRunSequenceCollection(repo, seq_cls=<class
                                                         'aim.sdk.sequence.Sequence'>,
                                                         query="", paginated=False,
                                                         offset=None, re-
                                                         port_mode=QueryReportMode.PROGRESS_BAR,
                                                         timezone_offset=0)
```

Bases: [SequenceCollection](#)

Implementation of SequenceCollection interface for repository's runs matching given query.

Method *iter()* returns Sequence iterator which yields Sequence for current run's all sequences. Method *iter_runs()* returns SequenceCollection iterator from repository's runs matching given query.

Parameters

- **repo** (Repo) – Aim repository object.
- **seq_cls** (type) – The collection's sequence class. Sequences not matching to `seq_cls.allowed_dtypes` will be skipped. *Sequence* by default, meaning all sequences will match.

- **query** (str, optional) – Query expression. If specified, method *iter_runs()* will skip runs not matching the query. If not, method *iter_run()* will return SequenceCollection iterator for all runs in repository.

AIM STORAGE

37.1 aim.storage.arrayview module

class aim.storage.arrayview.**ArrayView**

Array of homogeneous elements with sparse indices. Interface for working with array as a non-sparse array is available for cases when index values are not important.

first()

First index and value of the array.

first_idx()

First index of the array.

first_value()

First value of the array.

indices()

Return sparse indices iterator.

Yields

Array's next sparse index.

indices_list()

Get sparse indices as a list.

indices_numpy()

Get sparse indices as numpy array.

items()

Return items iterator.

Yields

Tuple of array's next sparse index and value.

keys()

Return sparse indices iterator.

Yields

Array's next sparse index.

last()

Last index and value of the array.

last_idx()

Last index of the array.

last_value()

Last value of the array.

sparse_list()

Get sparse indices and values as :obj:`list`s.

sparse_numpy()

Get sparse indices and values as numpy arrays.

tolist()

Convert to values list

values()

Return values iterator.

Yields

Array's next value.

values_list()

Get values as a list.

values_numpy()

Get values as numpy array.

ANONYMIZED TELEMETRY

We constantly seek to improve Aim for the community. Telemetry data helps us immensely by capturing anonymous usage analytics and statistics. The “anonymous usage” means that no user-specific data is tracked. Aim will generate a unique random ID, which is not tied to any user data (such as e-mail, UNIX username, etc.). You will be notified when you run `aim up`, `aim server` commands, as well as when running any SDK method which is tracked by Aim. User-defined data is **NOT** tracked. Only the fact of certain event has happened is tracked. Here is the complete list of events tracked by Aim SDK:

- Aim UI start (`aim up`)
- Aim Remote Tracking server start (`aim server`)
- Aim Watcher service start (`aim-watcher start`)
- New Repo initialization/re-initialization of existing Repo (`aim init`)
- New Run created
- Run resumed

38.1 Motivation

Aim uses segment’s analytics toolkit to collect basic info about the usage:

- Anonymized stripped-down basic usage analytics;
- Anonymized number of experiments and run. We constantly improve the storage and UI for performance in case of many experiments. This type of usage analytics helps us to stay on top of the performance problem.

38.2 How to opt out

You can turn telemetry off by running the following CLI command:

```
aim telemetry off
```

In order to turn telemetry back on, run

```
aim telemetry on
```

Note: Setting `AIM_UI_TELEMETRY_ENABLED` environment variable to `0` is obsolete. You’ll get a warning message if the variable is set, while telemetry is turned on.

CHANGELOG

39.1 3.17.2 Mar 28, 2023

- Fix explorer crashing issue caused by adding a `displayName` property in Grouping component (KaroMourad)

39.2 3.17.1 Mar 24, 2023

- Avoid explorer crashing when accessing empty chart values (KaroMourad)

39.3 3.17.0 Mar 24, 2023

39.3.1 3.17.0 Mar 24, 2023 - Enhancements

- Expose `run_name` and `run_hash` parameters to the `aim.sdk.adapters.pytorch_lightning.AimLogger` adapter (constd)
- Add navigation link to Experiment page from the Run page (roubkar)
- Add navigation to explorers from the Run page (roubkar)
- Implement Metrics Explorer v2 via Base Explorer (KaroMourad)
- Add Text Explorer to filter and compare text (roubkar)
- Add groundwork for the UI kit v2 for improved usability (arsengit)

39.3.2 3.17.0 Mar 24, 2023 - Fixes

- Add support for Path type to the `repo` attribute of the Run class (emekaokoli19)
- Add support for `jax>0.4.0` (n-gao)
- Add `-y` option to Aim CLI commands (emekaokoli19)
- Fix issue with toggling lines visibility during live update (roubkar)
- Fix the issue when HF model doesn't have `num_labels` attribute (mihran113)
- Fix table cell scrolling issue in the Texts tab of the Run page (roubkar)

39.4 3.16.2 Mar 3, 2023

- Add exception-free mode to Aim (alberttorosyan)
- Expose `capture_terminal_logs` argument for `aim.sdk.adapters` classes (mihran113)
- Handle inconsistency between Sequence data and metadata (alberttorosyan)

39.5 3.16.1 Feb 27, 2023

- Pin package version `alembic>=1.5.0` (justinyu)
- Fix segment `flush()` issue with no internet access (alberttorosyan)
- Fix the issue with an empty-illustrations styles on Base explorers (KaroMourad)
- Add 'join community' popup to the sidebar (KaroMourad)
- Use non-strict mode when logging HF model metadata (alberttorosyan)
- Add `set()` method implementation in `ProxyTree/SubtreeView` classes (alberttorosyan)

39.6 3.16.0 Feb 3, 2023

39.6.1 3.16.0 Feb 3, 2023 - Enhancements

- Drop support for python3.6 (mihran113)
- Add support for python3.11 (alberttorosyan)
- Add other x-axis alignment and system logs tracking to cli convert wandb (hjoonjang)
- Add support for pre-binned distribution/histogram (YodaEmbedding)
- Display logged run messages in Run page (VkoHov, alberttorosyan, roubkar)
- Use read-only mode when opening container for indexing (alberttorosyan)
- Add Stable-Baselines3 integration (tmynn)
- Add Acme integration (tmynn)
- Add huggingface/datasets integration (tmynn)
- Enable support for protobuf v4 (mihran113)
- Support events signaling for Remote Tracking server (alberttorosyan)
- Enhance DVC parameters tracking (tmynn)
- Add SDK events tracking (alberttorosyan)
- Add the ability to easily copy run hash with a single click (VkoHov)
- Add Prophet integration (grigoryan-davit)
- Add 'Dataset' type support for hf/datasets (tmynn)
- Add HuggingFace Transformers model info (tmynn)
- Add multidataset logging support for HuggingFace transformers (tmynn)

39.6.2 3.16.0 Feb 3, 2023 - Fixes

- Fix gpu stats logging when some stats are unavailable (timokau)
- Sub-path support for RTS addresses (mihran113)
- Fix experiment name update issues (mihran113)
- Fix run experiment setting race conditions (mihran113)
- Fix the issue with runs not appearing on UI (mihran113)

39.7 3.15.2 Dec 23, 2022

- Change logging level for reporter debug messages (alberttorosyan)
- Fix styling issues on the experiment page (KaroMourad)
- Fix client side worker port calculation for RTS (mihran113)
- Add discord community link in the sidebar (arsengit)
- Display experiments descriptions in the explorers tables (arsengit)

39.8 3.15.1 Dec 1, 2022

- Fix issue with index container lock for older repos (mihran113)
- Fix issue with rendering incorrect empty-illustration content in Audios explorer (KaroMourad)

39.9 3.15.0 Nov 26, 2022

39.9.1 3.15.0 Nov 26, 2022 - Enhancements:

- Implement Aim callbacks system and extended notifications (alberttorosyan)
- Add chart legends to the Metrics Explorer (KaroMourad)
- Implement vertically scalable version of Remote Tracking (mihran113, alberttorosyan)
- Add the ability to search, filter, and compare audio through Audios Explorer (VkoHov)
- Add epoch tracking for PyTorch Lightning (tmynn)
- Add PaddlePaddle integration (tmynn)
- Add Optuna integration (tmynn)
- Use `packaging` to parse version strings (jangop)
- Implement the experiment page for the overall experiment info view (VkoHov)
- Implement dynamic flushing mechanism for CheckIns based on the flag (mahnerak)
- Implement robust locking and indexing mechanism for Aim Runs (alberttorosyan)

39.9.2 3.15.0 Nov 26, 2022 - Fixes:

- Fix multiple progress bars handling for terminal logs capturing (mihran113)
- Handle resources when multiple Ctrl-Cs are pressed (alberttorosyan)
- Remove non unicode symbols from `aim up` command logs (mihran113)
- Fix “Show Table Diff” for list type elements in runs, params and scatters explorers (kumarshreshtha)
- Support non-Latin chars for encoding in Aim UI (roubkar)
- Make new CheckIns always override the expiry date, consistent to what is documented (mahnerak)

39.10 3.14.4 Nov 11, 2022

- Fix dropdowns’ selected options losses in time of searching other options in Figures Explorer (rubenaprikyan)
- Fix the group property name visibility in the images and audio tabs (VkoHov)
- Change the color contrast of the icons in the manage columns popover (VkoHov)
- Add notifier config files to aim package (alberttorosyan)
- Fix audios to numpy conversion (mihran113)

39.11 3.14.3 Oct 29, 2022

- Fix search for empty queries in explorers (KaroMourad)

39.12 3.14.2 Oct 28, 2022

- Add support to sync explorer state through url on Base and Figures Explorers (rubenaprikyan)
- Add support to highlight syntax error in Figures Explorer (KaroMourad)
- Fix issue with applying solid stroke styles on stroke badge in table (KaroMourad)
- Fix active runs indicators overlapping issue in LineChart (KaroMourad)
- Add support for text style formatting in the logs tab (VkoHov)
- Fix “TypeError: check() keywords must be strings” for `Run.metrics()` method (alberttorosyan)
- Fix run info API call error when tag color/description is None (alberttorosyan)
- Fix remote heartbeat resource cleanup (mihran113)

39.13 3.14.1 Oct 7, 2022

- Fix the current release duplication highlighting issue on the Dashboard page (arsengit)

39.14 3.14.0 Oct 6, 2022

39.14.1 3.14.0 Oct 6, 2022 - Enhancements:

- Move `aim reindex` command under `aim storage` group (mihran113)
- Add the ability to attach/remove tags on the Run Page (roubkar)
- Support dictionary as an argument of `Run.track` (alberttorosyan)
- Display the tags of the run in the tables of the explorers (VkoHov)
- Revamp Figures explorer controls and grouping sections for better onboarding and usability (VkoHov, KaroMourad)
- Replace the spinner loader with a lighter one (VkoHov)
- Add fast.ai integration (tmynn)
- Add command for dangling params cleanup (mihran113)
- Add top and bottom appearance modes to the chart popover (VkoHov)
- Deprecate Python 3.6 (alberttorosyan)
- Add MXNet integration (tmynn)
- Create a Dashboard page to provide a better onboarding experience (arsengit, roubkar, KaroMourad, mihran113)
- Add support for tracking jax device arrays (mihran113)

39.14.2 3.14.0 Oct 6, 2022 - Fixes:

- Fix chart hovering issue occurring when “nan” values are tracked (KaroMourad)
- Use empty dict as default when getting Run params (alberttorosyan)
- Change unit-tests data isolation mechanism (alberttorosyan)
- Adjust the visibility of the run color in tables (VkoHov)
- Fix response headers for remote tracking server (mihran113)
- Fix `TypeError`s in single run page (mihran113)

39.15 3.13.4 Sep 25, 2022

- Add the ability to disable smoothing explicitly (KaroMourad)
- Virtualize the run params list in the Run page (roubkar)

39.16 3.13.3 Sep 16, 2022

- Fix request cancellation on Logs tab (mihran113)
- Fix the data live update handling in the Logs tab (VkoHov)

39.17 3.13.2 Sep 10, 2022

- Fix content overlapping issue of x-axis alignment dropdown (KaroMourad)
- Fix the regression line rendering issue on Scatter plot exported image (KaroMourad)

39.18 3.13.1 Sep 1, 2022

- Add support for querying metrics by last value (mihran113)
- Fix aim reindex command failure (alberttorosyan)
- Fix issue with remote runs re-open (mihran113)
- Deprecate custom set Run.hash values (alberttorosyan)
- Tune mlflow converter run properties (tmynn)
- Fix AimLogger deprecation issues related to release of PyTorch Lightning v1.7 (djwessel)

39.19 3.13.0 Aug 21, 2022

39.19.1 3.13.0 Aug 21, 2022 - Enhancements:

- Add Figures Explorer to visualize and compare plotly figures (rubenaprikyan, KaroMourad, arsengit, VkoHov, roubkar)
- Add Base Explorer as core of all explorers (rubenaprikyan, KaroMourad, arsengit, VkoHov, roubkar)
- Add logging for remote resource cleanup and network stability (mihran113)
- Restrict Run.hash to auto-generated values only (alberttorosyan)
- Add ability to compare selected runs from the table (arsengit)
- Notify users about failed/stalled runs (mahnerak, alberttorosyan)
- Add ability to pin metrics in Run Page (mihran113, roubkar)
- Add step for unit tests for nightly releases workflow (mihran113)
- Add Keras-Tuner integration (tmynn)

- Add Weights & Biases to Aim log converter (tmynn)

39.19.2 3.13.0 Aug 21, 2022 - Fixes:

- Fix chart exporting issue (KaroMourad)
- Fix aim ui rendering issue on notebooks (rubenaprikyan)
- Fix live update retry to show live data after solving connection problems with the server (rubenaprikyan)
- Fix tensorboard convert while converting tensor (sharathmk99)
- Fix incorrect column keys of metrics in the table grid of the runs dashboard (VkoHov)
- Fix git info collection (mihran113)
- Fix code block content and query copying functionality (arsengit)
- Provide compatibility between plotly and matplotlib (tmynn)
- Warn to use aim.Image if aim.Figure fails (tmynn)

39.20 3.12.2 Aug 5, 2022

- Fix formatting of empty metric contexts (VkoHov)
- Apply lazy loading on metrics in Run Page (roubkar)

39.21 3.12.1 Aug 2, 2022

- Loosen version requirements for grpcio (alberttorosyan)
- Fix remote heartbeat-watcher resource cleanup (mihran113)
- Break long metric names into multiple lines in Run Page (roubkar)
- Enable run filtering by metric values (mihran113)
- Fix Cython version to eliminate build errors (mihran113)

39.22 3.12.0 Jul 22, 2022

39.22.1 3.12.0 Jul 22, 2022 - Enhancements:

- Add ability to set axes range manually for line charts on UI (KaroMourad)
- Add more user-friendly querying for dates (mihran113, arsengit)
- Filter redundant tooltip data from URL config state (KaroMourad)
- Improve rendering performance by enhancing table columns virtualization mechanism (roubkar)
- Increase visibility and usability of the Show table diff button (arsengit)
- Add support for tensorboard audios conversion (mihran113)
- Format params keys/paths properly (VkoHov)

- Mention explicitly run params everywhere params is mentioned (VkoHov)
- Add ability to hide a batch of items in explorers (VkoHov)
- Add ability to sort by the last value of the metric in table (VkoHov)
- Preserve active line even if it is dropped out of the filtered area (VkoHov)
- Add run duration property for SDK and queries (mihran113)
- Add client vs server version check for remote tracking server (mihran113)
- Add Remote tracking client heartbeat (mihran113)

39.22.2 3.12.0 Jul 22, 2022 - Fixes:

- Tune table sorting icon box overlapping with column box in compact mode (KaroMourad)
- Fix tensorboard log conversion for images (mihran113)
- Check if gradient is None when tracking gradient distributions (kage08)
- Fix displaying non-syntax errors across Aim UI (arsengit)
- Fix queries on remote repos (mihran113)
- Fix interval progress reports for query apis (mihran113)
- Fix query request cancellation errors (mihran113)
- Auto-detect and address inconsistencies in meta and series trees (mahnerak)

39.23 3.11.2 Jul 8, 2022

39.23.1 3.11.2 Jul 8, 2022 - Enhancements:

- Display the error position when getting syntax errors after searching (arsengit)

39.23.2 3.11.2 Jul 8, 2022 - Fixes:

- Avoid saving crashed or terminated search requests as the last state on explorers (arsengit)
- Remove the progress bar blinking when searching runs in Runs Explorer (KaroMourad)
- Fix the “matched runs” sentence color style in progress bars (KaroMourad)
- Fix `SyntaxError` handling for python3.10+ (mihran113)
- Fix generic Exceptions handling and adjust `HTTPException` handling (alberttorosyan)

39.24 3.11.1 Jun 27, 2022

- Replace base58 encoder with base64 (KaroMourad, VkoHov)
- Fix Notes tab loading issue (arsengit)
- Fix the loading logic of the `monaco editor` across the Aim Ui (arsengit)
- Fix Table export functionality in Params and Scatters explorers (arsengit)
- Allow mixing numeric types on a single Sequence (alberttorosyan)

39.25 3.11.0 Jun 21, 2022

39.25.1 3.11.0 Jun 21, 2022 - Enhancements:

- Add `--uds` option for `aim up` command (mihran113)
- Add progress reporting for search APIs and `tqdm` progress for SDK queries (mihran113)
- Add all the attributes of runs in the grouping popovers (KaroMourad)
- Display progress bar on Explorer pages when searching metadata (KaroMourad)
- Improve the processing speed for `tb to aim converter` (osoblanco)
- Adjust charts hover attributes position calculation and styles (KaroMourad)
- Improve formatting of numbers by setting maximum precision (KaroMourad)
- Add cloud storage backups to AWS S3 for aim repo runs (karan2801)
- Add LightGBM integration example (gorarakelyan)
- Add descriptive document titles for pages (KaroMourad)
- Implement unit-tests for aim SDK utils (yeghiakoronian)
- Display `std.dev/err` aggregated values in the table (VkoHov)
- Add `active` state indicator property for `aim.Run` (mihran113)
- Add `active` state indicators on the chart (VkoHov)
- Add ability to edit run name and description of run (VkoHov)
- Show the description in the sidebar of the run overview tab (VkoHov)
- Add all the attributes of run in the tooltip (VkoHov)
- Optimize the initial render time of Aim UI by using more lightweight font-family (arsengit)
- Use `monaco editor` as the syntax highlighter across the Aim UI (arsengit)
- Add loader to the top of the logs box in the run page (VkoHov)
- Add the date and the duration of run in the header of the single run page (VkoHov)
- Add the name, status and duration of run in the runs table of the tags page (VkoHov)
- Fit long name values in manage columns popover (arsengit)
- Add caching mechanism for sequence queries to optimize query performance (mihran113)
- Use step random hash as a key for metric sequences (alberttorosyan)

39.25.2 3.11.0 Jun 21, 2022 - Fixes:

- Fix issue with tensorboard to aim conversion (osoblanco)
- Fix reset zoom history on alignment type change (KaroMourad)
- Fix issue with rendering incorrect data when x-axis aligned by `relative time/epoch` (KaroMourad)
- Fix LineCart axis ticks overlapping issue on log scale (KaroMourad)
- Change zooming default option to multiple (VkoHov)
- Change grouped rows' min and max values names to `Group Min` and `Group Max` (VkoHov)
- Preserve the search input value of the grouping dropdown (VkoHov)
- Change the titles and placeholders in popovers (VkoHov)
- Resolve typing latency issue in the query search input (arsengit)
- Reorder and add non-hideable table columns (arsengit)
- Change the font of the runs navigation popover (VkoHov)
- Keep color persistence state after page reload (VkoHov)
- Resolve content blinking issue after search in the run page (arsengit)
- Fix scroll to bottom on live-update in logs tab (VkoHov)
- Fix timezone issues for activity map (mihran113)
- Fix `aim up` command output when `--port 0` is passed (mihran113)

39.26 3.10.3 May 31, 2022

- Adjust the content overflowing of the Delete and the Archive modals (VkoHov)
- Resolve issue with redirect in run page (arsengit)

39.27 3.10.2 May 26, 2022

- Adjust SRP Logs row height calculation (VkoHov)
- Fix issue with live update requests scheduler (rubenaprikyan)
- Fix log capturing crash during run garbage collection (mihran113)
- Fix Pytorch Lightning adapter `finalize` method (mihran113)
- Fix params duplication in dropdowns (VkoHov)
- Skip system params in Explorer pages (alberttorosyan)

39.28 3.10.1 May 18, 2022

- Resolve issue with rendering run params in the overview tab of SRP (arsengit)
- Fix issue with search query state update (arsengit)

39.29 3.10.0 May 17, 2022

39.29.1 3.10.0 May 17, 2022 - Enhancements:

- Add ability to adjust the density of the visible content in tables (roubkar)
- Set `metric.name` as default option for grouping (roubkar)
- Show user-selected params before group config in chart popover (roubkar)
- Optimize stream decoding performance on UI (mahnerak)
- Add support for animated image formats to Aim Image object (devfox-se)
- Add `AimLogger` for Catboost (devfox-se)
- Add `AimCallback` for LightGBM (devfox-se)
- Keep the extents of `HighPlot` axes brush in the state and the URL (VkoHov)
- Integrate aim with `cimport`-able aimrocks (mahnerak)
- Add `__slots__` to some classes to improve performance (mahnerak)
- Define base abstractions for `Iterator` and `DB` by borrowing from aimrocks (mahnerak)
- Use `KeysIterator` and `ValuesIterator` wrappers instead of reimplementing (mahnerak)
- Rename `PrefixView.container` to `PrefixView.parent` (mahnerak)
- Reimplement `absolute_path` (mahnerak)
- Cython bindings for `PrefixView`, `TreeView`, `Container`, `ArrayView` (mahnerak)
- Add ability to track and visualize `stdout/stderr` (mihran113, VkoHov)
- Fix `AimLogger` deprecation issues related to release of PyTorch Lightning v1.5 (arnauddhaene)
- Enable better autocomplete experience with monaco editor (arsengit)
- Pre-loading and caching necessary resources, add pre-loader animation to Aim UI (arsengit)

39.29.2 3.10.0 May 17, 2022 - Fixes:

- Remove hard-coded installation of pre-requirements (mahnerak)
- Remove duplicate code from `TreeView` and `Container` methods (mahnerak)
- Fix issue with filtering metrics values in single run page (KaroMourad)

39.30 3.9.4 May 12, 2022

- Fix run remote tracking queue cleanup (mihran113)
- Fix HF callback before training access (mihran113)
- Fix compatibility with Jinja 3.1 (devfox-se)

39.31 3.9.3 May 10, 2022

- Fix affecting stroke types after changing color persistence (KaroMourad)

39.32 3.9.2 Apr 29, 2022

- Move aim_ui package data to separate directory (devfox-se)

39.33 3.9.1 Apr 29, 2022

- Move aim_ui package data to separate directory (devfox-se)

39.34 3.9.0 Apr 29, 2022

39.34.1 3.9.0 Apr 29, 2022 - Enhancements:

- Add Notes Tab to single run page (arsengit)
- Add the run name to the batch delete and the batch archive modals (VkoHov)
- Increase the scalability of rendering lines in charts (KaroMourad)
- Increase live update requests delay to prevent performance issues (rubenaprikryan)
- Change font-family to monospace in the Table component (arsengit)
- Add info message for single value sliders (VkoHov)
- Add --log-level argument for aim up/server commands (mihran113)
- Add notes backend api interface (devfox-se)
- Fix type hints in Repo class (uduse)

39.34.2 3.9.0 Apr 29, 2022 - Fixes:

- Fix LineChart y-dimension margin calculation (KaroMourad)
- Fix HighPlot lines partially rendering issue (KaroMourad)
- Fix HighPlot axis ticks overlapping issue (KaroMourad)
- Fix sorting Params/Scatters explorer axis ticks (KaroMourad)
- Fix compatibility with pytorch-lightning v1.6.0 (mihran113)
- Fix the image's original size cropping (VkoHov)
- Fix PATH related issues for alembic and uvicorn (mihran113)
- Fix queries for custom object APIs (mihran113)
- Fix chart height updating when resize mode changed (VkoHov)
- Fix HuggingFace callback context capturing (mihran113)
- Fix Params/Scatters explorers' row hiding functionality (VkoHov)
- Fix Profiler logs are saved outside repo directory (devfox-se)

39.35 3.8.1 Apr 6, 2022

- Encode run hash before including in CSS selectors (Hamik25)
- Fix displaying incorrect metric values for large range scale in LineChart (KaroMourad)
- Fix issue with rendering lines for large range scale in LineChart (KaroMourad)
- Fix issue with URL state sync for bookmarks (roubkar)
- Fix issue with displaying negative param values on Aim UI (roubkar)
- Fix row hiding functionality (roubkar)
- Tune RunOverviewTab container styles (arsengit)
- Update documentations links on UI (rubenaprikyan)
- Fix RepoIndexManager run's reference cleanup (mihran113)
- Fix remote run finalization (mihran113)
- Fix issue with fetch on load more (infinite scroll) functionality in Runs Explorer (rubenaprikyan)

39.36 3.8.0 Mar 26, 2022**39.36.1 3.8.0 Mar 26, 2022 - Enhancements:**

- Hugging Face adapter refactoring (mihran113)
- Add run description columns to all run specific tables (VkoHov, mihran113)
- Change images rendering optimization default value to smoother (VkoHov)
- Set default steps ordering to desc in single run tabs (VkoHov, devfox-se)
- Add run name to grouping, ordering and run navigation popovers (VkoHov)

- Add ability to apply color scale on columns with numeric values (VkoHov)
- Refactored XGBoost AimCallback (devfox-se)
- Reopenable callbacks for integrations (mihran113)
- Add DVC integration (devfox-se)
- Add API profiler and unified API error response (devfox-se)
- Add API to retrieve N`th step of sequence (devfox-se)

39.36.2 3.8.0 Mar 26, 2022 - Fixes:

- Fix issue with calculation of active point on mouse hover in the LineChart (KaroMourad)
- Fix issue with wrong URL caching for Explorer pages (roubkar)
- Fix issue with focusing on the chart active point while moving the cursor (KaroMourad)
- Fix the image full view toggle icon visibility if the image has a white background (VkoHov)
- Fix scroll to the end of the audio tab (VkoHov)
- Add scrollbar to image full view mode content (VkoHov)
- Fix issues with run name/description not being set (mihran113)
- Fix issue with run single page tabs result caching (mihran113)
- Fix git system param tracking (devfox-se)
- Fix runs manual closing (mihran113)
- Fix Docker image creation step in packaging workflow (alberttorosyan)
- Fix Jinja2 template rendering with starlette==0.14.2 (alberttorosyan)

39.37 3.7.5 Mar 18, 2022

- Add request aborting functionality in single run page tabs (arsengit)
- Render plotly figures properly in single run page (arsengit)

39.38 3.7.4 Mar 15, 2022

- Fix density min and max validation calculation (VkoHov)

39.39 3.7.3 Mar 14, 2022

- Add missing names for dynamically imported files in single run page (arsengit)

39.40 3.7.2 Mar 10, 2022

- Fix issue with rendering UI re keeping long URL (KaroMourad)
- Split code in the single run page to optimize chunk size (arsengit)

39.41 3.7.1 Mar 10, 2022

- Fix metric queries with epoch=None (alberttorosyan)

39.42 3.7.0 Mar 9, 2022

39.42.1 3.7.0 Mar 9, 2022 - Enhancements:

- Add Run overview tab in run single page (arsengit, VkoHov, KaroMourad, rubenaprikyan)
- Custom max message size for Aim Remote tracking (alberttorosyan)
- Docker images for aim up/server (alberttorosyan)
- TF/Keras adapters refactoring (mihran113)
- Remote tracking client-side retry logic (aramaim)
- Add record_density to initial get-batch request for figures (VkoHov)

39.42.2 3.7.0 Mar 9, 2022 - Fixes:

- Fix rendering new lines in texts visualizer (arsengit)

39.43 3.6.3 Mar 4, 2022

- Fix UI rendering issue on colab (rubenaprikyan)

39.44 3.6.2 Mar 2, 2022

- Fix chart interactions issue in the Single Run Page Metrics tab (roubkar)
- Fix resolve_objects in remote tracking client subtree (alberttorosyan)
- Reject 0 as step/record count (alberttorosyan, VkoHov)
- Fix error on mlflow conversion by experiment id (devfox-se)

39.45 3.6.1 Feb 25, 2022

- Fix issue with aligning x-axis by custom metric (KaroMourad)
- Add `__AIM_PROXY_URL__` env variable to see full proxy url when running `aim up` command(rubenaprikyan)
- Add `--proxy-url` argument to notebook extension's `%aim up` to render UI correctly if there is a proxy server (rubenaprikyan)
- Add SageMaker integration, `jupyter-server-proxy` s bug-fix script (rubenaprikyan, mahnerak)
- Fix animation support in Plotly visualization and figure loading performance (Hamik25, mihran113)
- Display None values in group config column (VkoHov, Hamik25)
- Fix rendering issue on Select form search suggestions list (arsengit)
- Fix `PL.AimLogger save_dir` `AttributeError` (GeeekExplorer)
- Remove `__example_type__` substring from param name (VkoHov)

39.46 3.6.0 Feb 22 2022

39.46.1 3.6.0 Feb 22 2022 - Enhancements:

- Sort params columns in alphabetical order (arsengit)
- Add illustrations for indicating explorer search states (arsengit)
- Ability to export chart as image (KaroMourad)
- Ability to group by `metric.context` (VkoHov)
- Tune manage columns items highlighting styles (VkoHov)
- Set active style on table actions popover buttons with applied changes (arsengit)
- Unification of Run Custom Object APIs (alberttorosyan, VkoHov)
- Aim repo runs data automatic indexing (alberttorosyan)
- Pytorch Lightning adapter refactoring (mihran113)
- Add Pytorch Ignite integration (mihran113)
- Add wildcard support for `aim runs` subcommands (mihran113)
- Add MLflow logs conversion command (devfox-se)
- Add CustomObject implementation for `hub.dataset` (alberttorosyan)

39.46.2 3.6.0 Feb 22 2022 - Fixes:

- Fix live updated data loss after triggering endless scroll (VkoHov)
- Fix system metric columns pinning functionality and grouping column order (arsengit)
- Fix system metrics search in manage columns popover (VkoHov)
- Fix queries on remote repos (mihran113)
- Fix incorrect boolean value formatting (VkoHov)

39.47 3.5.4 Feb 15 2022

- Fix batch archive functionality (VkoHov)
- Add repo lock/release feature (devfox-se)

39.48 3.5.3 Feb 11 2022

- Fix rendering issue in runs explorer page (arsengit)

39.49 3.5.2 Feb 10 2022

- Fix issue with displaying current day activity cell on week's first day (rubenaprikyan)
- Fix issue with filtering options while typing in input of autocomplete in Tooltip and Grouping popovers (rubenaprikyan)

39.50 3.5.1 Feb 4 2022

- Fix folder creation when tracking with remote tracker (aramaim)

39.51 3.5.0 Feb 3 2022

39.51.1 3.5.0 Feb 3 2022 - Enhancements:

- Ability to hide system metrics from table (arsengit)
- Add input validations to range selectors (Hamik25)
- Improve media panel rendering performance on hovering over images (KaroMourad)
- Add ability to parse and import TensorFlow events into aim (devfox-se)
- Add system parameter logging: CLI, Env, Executable, Git, Installed packages (devfox-se)
- Convert nested non-native objects (e.g. OmegaConf config instance) upon storing (devfox-se)
- Add cli subcommands cp and mv for aim runs command (mihran113)
- Add handler for matplotlib figures in Image and Figure custom objects (devfox-se)

- Improve highlighting of table focused/hovered/selected row (VkoHov)

39.51.2 3.5.0 Feb 3 2022 - Fixes:

- Fix stalled runs deletion (mihran113)
- Fix background transparency in colab when using dark mode of system (rubenaprikyan)
- Fix Grouping and Tooltip popovers states' resetting issue when live-update is on (rubenaprikyan)
- Fix table column's sort functionality issue in Params and Scatters Explorers (rubenaprikyan)

39.52 3.4.1 Jan 23 2022

- Fix issue with displaying experiment name in Images Explorer table (VkoHov)

39.53 3.4.0 Jan 22 2022

- Add ability to apply group stacking on media elements list (KaroMourad)
- Add ability to apply sorting by run creation_time on table rows (roubkar)
- Add ability to filter texts table with keyword matching (roubkar, rubenaprikyan)
- Add ability to delete run from settings tab (Hamik25)
- Enhance controls states of explorer pages (arsengit)
- Add --repo, --host arguments support for notebook extension (VkoHov, rubenaprikyan)
- Add trendline options to ScatterPlot (roubkar)
- Add ability to display images in original size and align by width (arsengit)
- Add version, docs and slack links to sidebar (arsengit)
- Enhance AudioPlayer component (arsengit)
- Recover active tab in run details page after reload (roubkar)
- Add ability to archive or delete runs with batches (VkoHov)
- Remote tracking server [experimental] (alberttorosyan, mihran113, aramaim)
- Add ability to change media elements order (VkoHov)
- Add ability to hard delete runs (alberttorosyan)
- Lossy format support for aim.Image (devfox-se)
- Timezone issues fix for creation and end times (mihran113)

39.54 3.3.5 Jan 14 2022

- Add non-strict write mode to replace not-yet-supported types with their string representations. (mahnerak)
- Log pytorch_lightning hyperparameters in non-strict mode. (mahnerak)

39.55 3.3.4 Jan 10 2022

- Fix issue with WAL files flushing (alberttorosyan)
- Support for omegaconf configs in pytorch_lightning adapter (devfox-se)

39.56 3.3.3 Dec 24 2021

- Fix issue with showing range panel in Images Explorer (roubkar)

39.57 3.3.2 Dec 20 2021

- Fix issue with not providing point density value to live-update query (rubenaprikyan)

39.58 3.3.1 Dec 18 2021

- Fix getValue function to show correct chart title data (KaroMourad)

39.59 3.3.0 Dec 17 2021

- Add ability to track and explore audios in run detail page (arsengit, VkoHov, devfox-se)
- Add ability to track and visualize texts (mihran113, roubkar)
- Fix boolean values encoding (mahnerak)
- Add Scatter Explorer to visualize correlations between metric last value and hyperparameter (KaroMourad)
- Add ability to track and visualize plotly objects (devfox-se, Hamik25, rubenaprikyan)
- Add ability to query distributions by step range and density (VkoHov, rubenaprikyan)
- Add colab notebook support (mihran113, rubenaprikyan)
- Implement images visualization tab in run detail page (VkoHov, KaroMourad)
- Add custom URL prefix support (mihran113, Hamik25, roubkar)
- Enhance metric selection dropdowns to see lists in alphabetical order (rubenaprikyan)

39.60 3.2.2 Dec 10 2021

- Fix Run finalization index timeout issue (alberttorosyan)

39.61 3.2.1 Dec 8 2021

- Add ability to provide custom base path for API (mihran113, roubkar)
- Fix table groups column default order (arsengit)
- Fix table panel height issue in runs explorer page (arsengit)

39.62 3.2.0 Dec 3 2021

- Add ability to cancel pending request (roubkar, arsengit)
- Add support for secure protocol for API calls (mihran113, roubkar)
- Implement image full size view (VkoHov)
- Add ability to manipulate with image size and rendering type (arsengit)
- Enhance Table column for selected grouping config options (arsengit)
- Implement suggestions list for AimQL search (arsengit, rubenaprikyan)
- Add ability to track and visualize distributions (mihran113, rubenaprikyan)
- Add notebook extension, magic functions (rubenaprikyan)

39.63 3.1.1 Nov 25 2021

- Apply default ordering on images set (VkoHov)
- Ability to show image data in a tooltip on hover (KaroMourad)
- Support of Image input additional data sources (alberttorosyan)
- Ability to export run props as pandas dataframe (gorarakelyan)
- Slice image sequence by index for the given steps range (alberttorosyan)
- Improve Images Explorer rendering performance through better images list virtualization (roubkar)

39.64 3.1.0 Nov 20 2021

- Add ability to explore tracked images (VkoHov)
- Improve rendering performance by virtualizing table columns (roubkar)
- Add ability to apply grouping by higher level param key (roubkar)
- Add ability to specify repository path during `aim init` via `--repo` argument (rubenaprikyan)

39.65 3.0.7 Nov 17 2021

- Fix for missing metrics when numpy.float64 values tracked (alberttorosyan)

39.66 3.0.6 Nov 9 2021

- Fix for blocking container optimization for in progress runs (alberttorosyan)

39.67 3.0.5 Nov 9 2021

- Add tqdm package in setup.py required section (mihran113)

39.68 3.0.4 Nov 8 2021

- Switch to aimrocks 0.0.10 - exposes data flushing interface (mihran113)
- Optimize stored data when runs finalized (mihran113)
- Update aim reindex command to run storage optimizations (alberttorosyan)
- Storage partial optimizations on metric/run queries (alberttorosyan)

39.69 3.0.3 Nov 4 2021

- Bump sqlalchemy version to 1.4.1 (alberttorosyan)

39.70 3.0.2 Oct 27 2021

- Switch to aimrocks 0.0.9 - built on rocksdb 6.25.3 (alberttorosyan)
- Remove grouping select options from Params app config (VkoHov)
- Sort metrics data in ascending order for X-axis (KaroMourad)

39.71 3.0.1 Oct 22 2021

- Check telemetry_enabled option on segment initialization (VkoHov)
- Draw LineChart Y-axis (horizontal) tick lines on zooming (KaroMourad)
- Sort select options/params based on input value (roubkar)
- Fix query construction issue for multiple context items (roubkar)
- Fix issue with making API call from Web Worker (VkoHov)

39.72 3.0.0 Oct 21 2021

- Completely revamped UI:
 - Runs, metrics and params explorers
 - Bookmarks, Tags, Homepage
 - New UI works smooth with ~500 metrics displayed at the same time with full Aim table interactions
- Completely revamped storage:
 - 10x faster embedded storage based on Rocksdb
 - Average run query execution time on ~2000 runs: 0.784s
 - Average metrics query execution time on ~2000 runs with 6000 metrics: 1.552s

39.73 2.7.1 Jun 30 2021

- Fix bookmark navigation issue (roubkar)
- Empty metric select on X-axis alignment property change (roubkar)

39.74 2.7.0 Jun 23 2021

- Add ability to export table data as CSV (KaroMourad)
- Add ability to bookmark explore screen state (roubkar)
- Add dashboards and apps API (mihran113)

39.75 2.6.0 Jun 12 2021

- Resolve namedtuple python 3.5 incompatibility (gorarakelyan)
- Add ability to align X-axis by a metric (mihran113, roubkar)
- Add tooltip popover for the chart hover state (roubkar)

39.76 2.5.0 May 27 2021

- Set gunicorn timeouts (mihran113)
- Remove redundant deserialize method (gorarakelyan)
- Move the Flask server to main repo to support ‘docker’less UI (mihran113)

39.77 2.4.0 May 13 2021

- Bump up Aim UI to v1.6.0 (gorarakelyan)
- Add xgboost integration (khazhak)
- Update keras adapter interface (khazhak)
- Convert tensors to python numbers (gorarakelyan)

39.78 2.3.0 Apr 10 2021

- Bump up Aim UI to v1.5.0 (gorarakelyan)
- Set default interval of sys tracking to 10 seconds (gorarakelyan)
- Add ability to track system metrics (gorarakelyan)

39.79 2.2.1 Mar 31 2021

- Bump up Aim UI to v1.4.1 (gorarakelyan)

39.80 2.2.0 Mar 24 2021

- Bump up Aim UI to v1.4.0 (gorarakelyan)
- Add Hugging Face integration (Khazhak)
- Reorganize documentation (Tatevv)

39.81 2.1.6 Feb 26 2021

- Add ability to opt out telemetry (gorarakelyan)
- Remove experiment name from config file when calling repo.remove_branch method (gorarakelyan)

39.82 2.1.5 Jan 7 2021

- Handle NaN or infinite floats passed to artifacts (gorarakelyan)

39.83 2.1.4 Dec 2 2020

- Add ability to specify session run hash (gorarakelyan)
- Initialize repo if it was empty when opening session (gorarakelyan)
- Add validation of map artifact parameters (gorarakelyan)

39.84 2.1.3 Nov 24 2020

- Support comparison of list type contexts (gorarakelyan)

39.85 2.1.2 Nov 24 2020

- Fix empty contexts comparison issue (gorarakelyan)

39.86 2.1.1 Nov 22 2020

- Return only selected params in SelectResult (gorarakelyan)

39.87 2.1.0 Nov 19 2020

- Add AimRepo select method (gorarakelyan)
- Implement SelectResult class (gorarakelyan)

39.88 2.0.27 Nov 13 2020

- Fix issue with artifact step initializer (gorarakelyan)

39.89 2.0.26 Nov 10 2020

- Add `block_termination` argument to `aim.Session` (gorarakelyan)
- Convert infinity parameter to string in artifacts (gorarakelyan)

39.90 2.0.25 Nov 9 2020

- Reconstruct run metadata file when running close command (gorarakelyan)

39.91 2.0.24 Nov 8 2020

- Add SIGTERM signal handler (gorarakelyan)
- Run `track` function in a parallel thread (gorarakelyan)
- Add SDK session flush method (gorarakelyan)
- Flush aggregated metrics at a given frequency (gorarakelyan)
- Update run metadata file only on artifacts update (gorarakelyan)

39.92 2.0.23 Nov 5 2020

- Make experiment name argument required in SDK close command (gorarakelyan)

39.93 2.0.22 Nov 5 2020

- Add SDK `close` method to close dangling experiments (gorarakelyan)

39.94 2.0.21 Nov 1 2020

- Resolve compatibility issues with python 3.5.0 (gorarakelyan)

39.95 2.0.20 Oct 26 2020

- Enable pypi aim package name (gorarakelyan)

39.96 2.0.19 Oct 25 2020

- Add PyTorch Lightning logger (gorarakelyan)
- Add TensorFlow v1 and v2 keras callbacks support (gorarakelyan)

39.97 2.0.18 Oct 7 2020

- Add ability to run Aim UI in detached mode (gorarakelyan)
- Add ability to specify repo path when running Aim UI (gorarakelyan)

39.98 2.0.17 Oct 5 2020

- Rename AimDE to Aim UI (gorarakelyan)

39.99 2.0.16 Oct 2 2020

- Add ability to specify host when running AimDE (gorarakelyan)
- Disable AimContainerCommandManager (gorarakelyan)
- Remove aimde command entry point (gorarakelyan)
- Remove de prefix from development environment management commands (gorarakelyan)

39.100 2.0.15 Sep 21 2020

- Set Map artifact default namespace (gorarakelyan)

39.101 2.0.14 Sep 21 2020

- Set Metric hashable context to None if no kwarg is passed (gorarakelyan)

39.102 2.0.13 Sep 21 2020

- Add ability to query runs by metric value (gorarakelyan)
- Add ability to query runs via SDK (gorarakelyan)

39.103 2.0.12 Sep 12 2020

- Update Session to handle exceptions gracefully (gorarakelyan)

39.104 2.0.11 Sep 11 2020

- Add alias to keras adapter (gorarakelyan)

39.105 2.0.10 Sep 10 2020

- Show progress bar when pulling AimDE image (gorarakelyan)

39.106 2.0.9 Sep 10 2020

- Add ability to start multiple sessions (gorarakelyan)
- Add Aim adapter for keras (gorarakelyan)

39.107 2.0.8 Aug 26 2020

- Set SDK to select only unarchived runs by default (gorarakelyan)
- Add ability to archive/unarchive runs (gorarakelyan)
- Enable search by run attributes (gorarakelyan)
- Add `is not` keyword to AimQL (gorarakelyan)

39.108 2.0.7 Aug 21 2020

- Validate Artifact values before storing (gorarakelyan)
- Add sessions to SDK (gorarakelyan)

39.109 2.0.6 Aug 13 2020

- Add ability to retrieve metrics and traces from repo (gorarakelyan)
- Add SDK `select` method to select runs and artifacts (gorarakelyan)
- Implement search query language (gorarakelyan)

39.110 2.0.5 Jul 18 2020

- Fix issue with PyPI reStructuredText format compatibility (gorarakelyan)

39.111 2.0.4 Jul 18 2020

- Add ability to attach tf.summary logs to AimDE (gorarakelyan)

39.112 2.0.3 Jul 8 2020

- Pass project path to development environment container (gorarakelyan)

39.113 2.0.2 Jul 7 2020

- Make epoch argument optional for `Metric` artifact (gorarakelyan)
- Add ability to automatically commit runs after exit (gorarakelyan)
- Add `aim up` shortcut for running development environment (gorarakelyan)
- Remove first required argument (artifact name) from `sdk track` function (gorarakelyan)
- Add general dictionary artifact for tracking key: value parameters (gorarakelyan)

39.114 2.0.1 Jun 24 2020

- Fix inconsistent DE naming (gorarakelyan)

39.115 2.0.0 Jun 18 2020

- Tidy up aim and remove some artifacts (gorarakelyan)
- Update `AimContainerCMD` to open connection on custom port (gorarakelyan)
- Save passed process uuid to commit configs (gorarakelyan)
- Ability to query processes (gorarakelyan)
- Execute process and store logs into a commit of specific experiment (gorarakelyan)
- Kill running process and its children recursively (gorarakelyan)
- Keep executed processes for monitoring and management (gorarakelyan)
- Add container command handler to exec commands on the host (gorarakelyan)
- Refactor `Text` artifact to store sentences using `protobuf` and `aimrecords` (jamesj-jiao)
- Add ability to pass aim board port as an argument (gorarakelyan)

39.116 1.2.17 May 8 2020

- Add config command (gorarakelyan)
- Tune artifacts: images, metric_groups, params (gorarakelyan)

39.117 1.2.16 Apr 29 2020

- Add ability to pass numpy array as a segmentation mask (gorarakelyan)

39.118 1.2.15 Apr 29 2020

- Add basic image list tracking (gorarakelyan)

39.119 1.2.14 Apr 27 2020

- Optimize segmentation tracking insight to load faster (gorarakelyan)

39.120 1.2.13 Apr 25 2020

- Remove GitHub security alert (gorarakelyan)
- Add image semantic segmentation tracking (gorarakelyan)

39.121 1.2.12 Apr 20 2020

- Add missing init file for aim.artifacts.proto (@mike1808)

39.122 1.2.11 Apr 16 2020

- Make epoch property optional for Metric (gorarakelyan)

39.123 1.2.10 Apr 16 2020

- Serialize and store Metric records using protobuf and aimrecords (gorarakelyan)
- Create RecordWriter factory which handles artifact records saving (gorarakelyan)
- Extract artifact serialization to ArtifactWriter (mike1808)

39.124 1.2.9 Mar 16 2020

- Alert prerequisites installation message for running board (gorarakelyan)

39.125 1.2.8 Mar 15 2020

- Update profiler interface for keras (gorarakelyan)

39.126 1.2.7 Mar 14 2020

- Add board pull command (gorarakelyan)
- Change board ports to 43800,1,2 (gorarakelyan)
- Add ability to profile graph output nodes (gorarakelyan)
- Remove issue with autograd inside while loop (gorarakelyan)
- Add aim board development mode (gorarakelyan)
- Update board name hash algorithm to md5 (gorarakelyan)
- Add board CLI commands: up, down and upgrade (gorarakelyan)
- Add ability to tag version as a release candidate (gorarakelyan)

39.127 1.2.6 Feb 28 2020

- Add learning rate update tracking (gorarakelyan)

39.128 1.2.5 Feb 25 2020

- Add autocommit feature to push command: `aim push -c [-m <msg>]` (gorarakelyan)
- Add cli status command to list branch uncommitted artifacts (gorarakelyan)
- Add an ability to aggregate duplicated nodes within a loop (gorarakelyan)
- Remove gradient break issue when profiling output nodes (gorarakelyan)

39.129 1.2.4 Feb 20 2020

- Enable profiler to track nodes inside loops (gorarakelyan)
- Ability to disable profiler for evaluation or inference (gorarakelyan)

39.130 1.2.3 Feb 13 2020

- Set minimum required python version to 3.5.2 (gorarakelyan)

39.131 1.2.2 Feb 13 2020

- Downgrade required python version (gorarakelyan)

39.132 1.2.1 Feb 13 2020

- Edit README.md to pass reStructuredText validation on pypi (gorarakelyan)

39.133 1.2.0 Feb 13 2020

- Make aim CLI directly accessible from main.py (gorarakelyan)
- Add disk space usage tracking (gorarakelyan)
- Add profiler support for Keras (gorarakelyan)
- Add TensorFlow graph nodes profiler (gorarakelyan)
- Add command to run aim live container mounted on aim repo (gorarakelyan)
- Update profiler to track GPU usage (gorarakelyan)
- Add machine resource usage profiler (gorarakelyan)

39.134 1.1.1 Jan 14 2020

- Remove aim dependencies such as keras, pytorch and etc (gorarakelyan)

39.135 1.1.0 Jan 12 2020

- Update code diff tracking to be optional (gorarakelyan)
- Add default False value to aim init function (gorarakelyan)
- Update aim repo to correctly identify cwd (gorarakelyan)
- Update push command to commit if msg argument is specified (gorarakelyan)
- Add ability to initialize repo from within the sdk (gorarakelyan)

39.136 1.0.2 Jan 7 2020

- Remove objects dir from empty .aim branch index (gorarakelyan)

39.137 1.0.1 Dec 26 2019

- Add cil command to print aim current version (gorarakelyan)

39.138 1.0.0 Dec 25 2019

- Add aim version number in commit config file (gorarakelyan)
- Update push command to send username and check storage availability (gorarakelyan)
- Add hyper parameters tracking (gorarakelyan)
- Update push command to print shorter file names when pushing to remote (gorarakelyan)
- Update tracking artifacts to be saved in log format (gorarakelyan)
- Add pytorch cuda support to existing sdk artefacts (gorarakelyan)
- Add cli reset command (gorarakelyan)
- Add nested module tracking support to aim sdk (gorarakelyan)
- Add code difference tracking to aim sdk (gorarakelyan)
- Update aim push command to send commits (gorarakelyan)
- Add commit structure implementation (gorarakelyan)
- Add aim commit command synchronized with git commits (gorarakelyan)
- Add version control system factory (gorarakelyan)
- Update all insights example (gorarakelyan)
- Add model gradients tracking (gorarakelyan)
- Add model weights distribution tracking (gorarakelyan)
- Add aim correlation tracking (gorarakelyan)

39.139 0.2.9 Nov 30 2019

- Update push tolerance when remote origin is invalid (gorarakelyan)

39.140 0.2.8 Nov 30 2019

- Update aim auth public key search algorithm (gorarakelyan)

39.141 0.2.7 Nov 14 2019

- Update dependencies torch and torchvision versions (sgevorg)

39.142 0.2.6 Nov 5 2019

- Update aim track logger (gorarakelyan)

39.143 0.2.5 Nov 4 2019

- Add branch name validation (gorarakelyan)
- Add single branch push to aim push command (gorarakelyan)

39.144 0.2.4 Nov 3 2019

- Update aim auth print format (gorarakelyan)
- Update setup.py requirements (gorarakelyan)

39.145 0.2.3 Nov 3 2019

- Update package requirements (gorarakelyan)

39.146 0.2.2 Nov 1 2019

- Update package requirements (sgevorg)

39.147 0.2.1 Nov 1 2019

- Add paramiko to required in setup.py (sgevorg)

39.148 0.2.0 Nov 1 2019

- Update the repo to prep for open source pypi push (sgevorg)
- Add error and activity logging (sgevorg)
- Add push command robustness (gorarakelyan)
- Add cli auth command (gorarakelyan)
- Add public key authentication (gorarakelyan)
- Update push to send only branches (gorarakelyan)
- Add branching command line interface (gorarakelyan)
- Update skd interface (gorarakelyan)
- Add pytorch examples inside examples directory (gorarakelyan)
- Add model load sdk method (gorarakelyan)
- Add model checkpoint save tests (gorarakelyan)
- Update file sending protocol (gorarakelyan)
- Add model tracking (gorarakelyan)

39.149 0.1.0 - Sep 23 2019

- Update setup py to build cython extensions (gorarakelyan)
- Update tcp client to send multiple files through one connection (gorarakelyan)
- Update tcp client to send images (gorarakelyan)
- Update sdk track functionality to support multiple metrics (gorarakelyan)
- Update push command for sending repo to a given remote (gorarakelyan)
- Add cli remote commands (gorarakelyan)
- Update cli architecture from single group of commands to multiple groups (gorarakelyan)
- Add testing env first skeleton and versions (sgevorg)
- Add dummy exporting files from .aim-test (sgevorg)
- Add description for Testing Environment (sgevorg)
- Update metadata structure and handling (sgevorg)
- Add support for seq2seq models (sgevorg)
- Update the output of docker image build to be more informative and intuitive (sgevorg)
- Update README.MD with changed Aim messaging (sgevorg)
- Remove setup.cfg file (maybe temporarily) (sgevorg)
- Update the location for docker build template files, move to data/ (sgevorg)
- Update the docs/cli.md for aim-deploy docs (sgevorg)
- Add docker deploy .aim/deploy_temp/<model> cleanup at the end of the build (sgevorg)
- Add Docker Deploy via aim-deploy command (sgevorg)

- Add Docker image generate skeleton (sgevorg)
- Add AimModel.load_mode static function to parse .aim files (sgevorg)
- Update exporter to decouple from specifics of exporting and framework (sgevorg)
- Add model export with .aim extension (sgevorg)
- Remove pack/unpack of the metadata (sgevorg)
- Add pack/unpack to add metadata to model for engine processing (sgevorg)
- Add aim-deploy command configuration in cli (sgevorg)
- Add basic cli (sgevorg)
- Update setup.py for cli first version (sgevorg)
- Add initial cli specs (sgevorg)
- Add directories: the initial skeleton of the repo (sgevorg)
- Add gitignore, license file and other basics for repo (sgevorg)

INDICES AND TABLES

- `genindex`
- `modindex`

PYTHON MODULE INDEX

a

- `aim.sdk.objects.audio`, 147
- `aim.sdk.objects.distribution`, 146
- `aim.sdk.objects.figure`, 148
- `aim.sdk.objects.image`, 145
- `aim.sdk.objects.text`, 148
- `aim.sdk.repo`, 135
- `aim.sdk.run`, 139
- `aim.sdk.sequence`, 148
- `aim.sdk.sequence_collection`, 150
- `aim.sdk.sequences.audio_sequence`, 150
- `aim.sdk.sequences.distribution_sequence`, 149
- `aim.sdk.sequences.figure_sequence`, 150
- `aim.sdk.sequences.image_sequence`, 149
- `aim.sdk.sequences.text_sequence`, 150
- `aim.storage.arrayview`, 153

Symbols

`__delitem__()` (*aim.sdk.run.Run* method), 139
`__getitem__()` (*aim.sdk.run.Run* method), 140
`__setitem__()` (*aim.sdk.run.Run* method), 140

A

`active` (*aim.sdk.run.Run* property), 143
`add_tag()` (*aim.sdk.run.Run* method), 140
`aim.sdk.objects.audio`
 module, 147
`aim.sdk.objects.distribution`
 module, 146
`aim.sdk.objects.figure`
 module, 148
`aim.sdk.objects.image`
 module, 145
`aim.sdk.objects.text`
 module, 148
`aim.sdk.repo`
 module, 135
`aim.sdk.run`
 module, 139
`aim.sdk.sequence`
 module, 148
`aim.sdk.sequence_collection`
 module, 150
`aim.sdk.sequences.audio_sequence`
 module, 150
`aim.sdk.sequences.distribution_sequence`
 module, 149
`aim.sdk.sequences.figure_sequence`
 module, 150
`aim.sdk.sequences.image_sequence`
 module, 149
`aim.sdk.sequences.text_sequence`
 module, 150
`aim.storage.arrayview`
 module, 153
`allowed_dtypes()` (*aim.sdk.sequence.Sequence* class method), 148
`archived` (*aim.sdk.run.Run* property), 143
`ArrayView` (class in *aim.storage.arrayview*), 153

`Audio` (class in *aim.sdk.objects.audio*), 147
`Audios` (class in *aim.sdk.sequences.audio_sequence*), 150

B

`bin_count` (*aim.sdk.objects.distribution.Distribution* property), 147

C

`caption` (*aim.sdk.objects.image.Image* property), 145
`collect_params_info()` (*aim.sdk.repo.Repo* method), 135
`collect_sequence_info()` (*aim.sdk.repo.Repo* method), 135
`collect_sequence_info()` (*aim.sdk.run.Run* method), 140
`copy_runs()` (*aim.sdk.repo.Repo* method), 135
`creation_time` (*aim.sdk.run.Run* property), 144

D

`dataframe()` (*aim.sdk.run.Run* method), 140
`default_repo()` (*aim.sdk.repo.Repo* class method), 136
`delete_run()` (*aim.sdk.repo.Repo* method), 136
`delete_runs()` (*aim.sdk.repo.Repo* method), 136
`description` (*aim.sdk.run.Run* property), 144
`Distribution` (class in *aim.sdk.objects.distribution*), 146
`Distributions` (class in *aim.sdk.sequences.distribution_sequence*), 149
`duration` (*aim.sdk.run.Run* property), 144

E

`end_time` (*aim.sdk.run.Run* property), 144
`epochs` (*aim.sdk.sequence.Sequence* property), 149
`exists()` (*aim.sdk.repo.Repo* class method), 136
`experiment` (*aim.sdk.run.Run* property), 144

F

`Figure` (class in *aim.sdk.objects.figure*), 148
`Figures` (class in *aim.sdk.sequences.figure_sequence*), 150

`first()` (*aim.storage.arrayview.ArrayView method*), 153
`first_idx()` (*aim.storage.arrayview.ArrayView method*), 153
`first_step()` (*aim.sdk.sequence.Sequence method*), 148
`first_value()` (*aim.storage.arrayview.ArrayView method*), 153
`format` (*aim.sdk.objects.image.Image property*), 145
`from_histogram()` (*aim.sdk.objects.distribution.Distribution class method*), 146
`from_path()` (*aim.sdk.repo.Repo class method*), 136
`from_samples()` (*aim.sdk.objects.distribution.Distribution class method*), 146

G

`get()` (*aim.sdk.objects.audio.Audio method*), 148
`get_audio_sequence()` (*aim.sdk.run.Run method*), 141
`get_distribution_sequence()` (*aim.sdk.run.Run method*), 141
`get_figure_sequence()` (*aim.sdk.run.Run method*), 141
`get_image_sequence()` (*aim.sdk.run.Run method*), 141
`get_log_records()` (*aim.sdk.run.Run method*), 141
`get_metric()` (*aim.sdk.run.Run method*), 141
`get_run()` (*aim.sdk.repo.Repo method*), 137
`get_terminal_logs()` (*aim.sdk.run.Run method*), 142
`get_text_sequence()` (*aim.sdk.run.Run method*), 142

H

`height` (*aim.sdk.objects.image.Image property*), 146

I

`Image` (*class in aim.sdk.objects.image*), 145
`Images` (*class in aim.sdk.sequences.image_sequence*), 149
`indices()` (*aim.storage.arrayview.ArrayView method*), 153
`indices_list()` (*aim.storage.arrayview.ArrayView method*), 153
`indices_numpy()` (*aim.storage.arrayview.ArrayView method*), 153
`items()` (*aim.storage.arrayview.ArrayView method*), 153
`iter()` (*aim.sdk.sequence_collection.SequenceCollection method*), 150
`iter_metrics_info()` (*aim.sdk.run.Run method*), 142
`iter_runs()` (*aim.sdk.repo.Repo method*), 137
`iter_runs()` (*aim.sdk.sequence_collection.SequenceCollection method*), 150
`iter_sequence_info_by_type()` (*aim.sdk.run.Run method*), 142

J

`json()` (*aim.sdk.objects.audio.Audio method*), 148
`json()` (*aim.sdk.objects.distribution.Distribution method*), 147
`json()` (*aim.sdk.objects.figure.Figure method*), 148
`json()` (*aim.sdk.objects.image.Image method*), 145

K

`keys()` (*aim.storage.arrayview.ArrayView method*), 153

L

`last()` (*aim.storage.arrayview.ArrayView method*), 153
`last_idx()` (*aim.storage.arrayview.ArrayView method*), 153
`last_step()` (*aim.sdk.sequence.Sequence method*), 149
`last_value()` (*aim.storage.arrayview.ArrayView method*), 154

M

`metrics()` (*aim.sdk.run.Run method*), 142
`module`
 aim.sdk.objects.audio, 147
 aim.sdk.objects.distribution, 146
 aim.sdk.objects.figure, 148
 aim.sdk.objects.image, 145
 aim.sdk.objects.text, 148
 aim.sdk.repo, 135
 aim.sdk.run, 139
 aim.sdk.sequence, 148
 aim.sdk.sequence_collection, 150
 aim.sdk.sequences.audio_sequence, 150
 aim.sdk.sequences.distribution_sequence, 149
 aim.sdk.sequences.figure_sequence, 150
 aim.sdk.sequences.image_sequence, 149
 aim.sdk.sequences.text_sequence, 150
 aim.storage.arrayview, 153
`move_runs()` (*aim.sdk.repo.Repo method*), 137

N

`name` (*aim.sdk.run.Run property*), 144

Q

`query_audios()` (*aim.sdk.repo.Repo method*), 137
`query_distributions()` (*aim.sdk.repo.Repo method*), 137
`query_figure_objects()` (*aim.sdk.repo.Repo method*), 137
`query_images()` (*aim.sdk.repo.Repo method*), 138
`query_metrics()` (*aim.sdk.repo.Repo method*), 138
`query_runs()` (*aim.sdk.repo.Repo method*), 138
`query_texts()` (*aim.sdk.repo.Repo method*), 139

QueryRunSequenceCollection (class in *aim.sdk.sequence_collection*), 151

QuerySequenceCollection (class in *aim.sdk.sequence_collection*), 151

R

range (*aim.sdk.objects.distribution.Distribution* property), 147

ranges (*aim.sdk.objects.distribution.Distribution* property), 147

remove_tag() (*aim.sdk.run.Run* method), 143

Repo (class in *aim.sdk.repo*), 135

report_progress() (*aim.sdk.run.Run* method), 143

report_successful_finish() (*aim.sdk.run.Run* method), 143

rm() (*aim.sdk.repo.Repo* class method), 139

Run (class in *aim.sdk.run*), 139

S

Sequence (class in *aim.sdk.sequence*), 148

sequence_name() (*aim.sdk.sequence.Sequence* class method), 149

sequence_name() (*aim.sdk.sequences.audio_sequence.Audios* class method), 150

sequence_name() (*aim.sdk.sequences.distribution_sequence.Distributions* class method), 149

sequence_name() (*aim.sdk.sequences.figure_sequence.Figures* class method), 150

sequence_name() (*aim.sdk.sequences.image_sequence.Images* class method), 149

sequence_name() (*aim.sdk.sequences.text_sequence.Texts* class method), 150

SequenceCollection (class in *aim.sdk.sequence_collection*), 150

SingleRunSequenceCollection (class in *aim.sdk.sequence_collection*), 150

size (*aim.sdk.objects.image.Image* property), 146

sparse_list() (*aim.storage.arrayview.ArrayView* method), 154

sparse_numpy() (*aim.storage.arrayview.ArrayView* method), 154

T

tags (*aim.sdk.run.Run* property), 144

Text (class in *aim.sdk.objects.text*), 148

Texts (class in *aim.sdk.sequences.text_sequence*), 150

timestamps (*aim.sdk.sequence.Sequence* property), 149

to_np_histogram() (*aim.sdk.objects.distribution.Distribution* method), 147

to_numpy() (*aim.sdk.objects.audio.Audio* method), 148

to_pil_image() (*aim.sdk.objects.image.Image* method), 145

tolist() (*aim.storage.arrayview.ArrayView* method), 154

track() (*aim.sdk.run.Run* method), 143

V

values (*aim.sdk.sequence.Sequence* property), 149

values() (*aim.storage.arrayview.ArrayView* method), 154

values_list() (*aim.storage.arrayview.ArrayView* method), 154

values_numpy() (*aim.storage.arrayview.ArrayView* method), 154

W

weights (*aim.sdk.objects.distribution.Distribution* property), 147

width (*aim.sdk.objects.image.Image* property), 146