

---

**hi-ml**

**InnerEye**

**Oct 20, 2021**



## WORKING WITH AZURE

<b>1</b>	<b>First steps: How to run your Python code in the cloud</b>	<b>3</b>
<b>2</b>	<b>Connecting to Azure</b>	<b>7</b>
<b>3</b>	<b>Datasets</b>	<b>9</b>
<b>4</b>	<b>Hyperparameter Search via Hyperdrive</b>	<b>13</b>
<b>5</b>	<b>Using Cheap Low Priority VMs</b>	<b>15</b>
<b>6</b>	<b>Commandline tools</b>	<b>19</b>
<b>7</b>	<b>Downloading from/ uploading to Azure ML</b>	<b>21</b>
<b>8</b>	<b>Examples</b>	<b>23</b>
<b>9</b>	<b>Logging metrics when training models in AzureML</b>	<b>29</b>
<b>10</b>	<b>Notes for developers</b>	<b>31</b>
<b>11</b>	<b>Contributing to this toolbox</b>	<b>35</b>
<b>12</b>	<b>health_azure Package</b>	<b>37</b>
<b>13</b>	<b>health_ml.utils Package</b>	<b>49</b>
<b>14</b>	<b>Indices and tables</b>	<b>53</b>
	<b>Python Module Index</b>	<b>55</b>
	<b>Index</b>	<b>57</b>



This toolbox helps to simplify and streamline work on deep learning models for healthcare and life sciences, by providing tested components (data loaders, pre-processing), deep learning models, and cloud integration tools.

The *hi-ml* toolbox provides

- Functionality to easily run Python code in Azure Machine Learning services
- Low-level and high-level building blocks for Machine Learning / AI researchers and practitioners.



## FIRST STEPS: HOW TO RUN YOUR PYTHON CODE IN THE CLOUD

The simplest use case for the `hi-ml` toolbox is taking a script that you developed, and run it inside of Azure Machine Learning (AML) services. This can be helpful because the cloud gives you access to massive GPU resource, you can consume vast datasets, and access multiple machines at the same time for distributed training.

### 1.1 Setting up AzureML

You need to have an AzureML workspace in your Azure subscription. Download the config file from your AzureML workspace, as described [here](#). **Put this file (it should be called `config.json`) into the folder where your script lives, or one of its parent folders.** You can use parent folders up to the last parent that is still included in the `PYTHONPATH` environment variable: `hi-ml` will try to be smart and search through all folders that it thinks belong to your current project.

### 1.2 Using the AzureML integration layer

Consider a simple use case, where you have a Python script that does something - this could be training a model, or pre-processing some data. The `hi-ml` package can help easily run that on Azure Machine Learning (AML) services.

Here is an example script that reads images from a folder, resizes and saves them to an output folder:

```
from pathlib import Path
if __name__ == '__main__':
    input_folder = Path("/tmp/my_dataset")
    output_folder = Path("/tmp/my_output")
    for file in input_folder.glob("*.jpg"):
        contents = read_image(file)
        resized = contents.resize(0.5)
        write_image(output_folder / file.name)
```

Doing that at scale can take a long time. **We'd like to run that script in AzureML, consume the data from a folder in blob storage, and write the results back to blob storage**, so that we can later use it as an input for model training.

You can achieve that by adding a call to `submit_to_azure_if_needed` from the `hi-ml` package:

```
from pathlib import Path
from health_azure import submit_to_azure_if_needed
if __name__ == '__main__':
    current_file = Path(__file__)
    run_info = submit_to_azure_if_needed(compute_cluster_name="preprocess-ds12",
```

(continues on next page)

(continued from previous page)

```

input_datasets=["images123"],
# Omit this line if you don't create an output_
↳ dataset (for example, in
# model training scripts)
output_datasets=["images123_resized"],
default_datastore="my_datastore")

# When running in AzureML, run_info.input_datasets and run_info.output_datasets will_
↳ be populated,
# and point to the data coming from blob storage. For runs outside AML, the paths_
↳ will be None.
# Replace the None with a meaningful path, so that we can still run the script_
↳ easily outside AML.
input_dataset = run_info.input_datasets[0] or Path("/tmp/my_dataset")
output_dataset = run_info.output_datasets[0] or Path("/tmp/my_output")
files_processed = []
for file in input_dataset.glob("*.jpg"):
    contents = read_image(file)
    resized = contents.resize(0.5)
    write_image(output_dataset / file.name)
    files_processed.append(file.name)

# Any other files that you would not consider an "output dataset", like metrics, etc,
↳ should be written to
# a folder "./outputs". Any files written into that folder will later be visible in_
↳ the AzureML UI.
# run_info.output_folder already points to the correct folder.
stats_file = run_info.output_folder / "processed_files.txt"
stats_file.write_text("\n".join(files_processed))

```

Once these changes are in place, you can submit the script to AzureML by supplying an additional `--azureml` flag on the commandline, like `python myscript.py --azureml`.

Note that you do not need to modify the argument parser of your script to recognize the `--azureml` flag.

### 1.3 Essential arguments to submit\_to\_azure\_if\_needed

When calling `submit_to_azure_if_needed`, you can supply the following parameters:

- `compute_cluster_name` (**Mandatory**): The name of the AzureML cluster that should run the job. This can be a cluster with CPU or GPU machines. See [here for documentation](#)
- `entry_script`: The script that should be run. If omitted, the `hi-ml` package will assume that you would like to submit the script that is presently running, given in `sys.argv[0]`.
- `snapshot_root_directory`: The directory that contains all code that should be packaged and sent to AzureML. All Python code that the script uses must be copied over. This defaults to the current working directory, but can be one of its parents. If you would like to explicitly skip some folders inside the `snapshot_root_directory`, then use `ignored_folders` to specify those.
- `conda_environment_file`: The conda configuration file that describes which packages are necessary for your script to run. If omitted, the `hi-ml` package searches for a file called `environment.yml` in the current folder or its parents.

You can also supply an input dataset. For data pre-processing scripts, you can add an output dataset (omit this for ML training scripts).



- To use datasets, you need to provision a data store in your AML workspace, that points to your training data in blob storage. This is described [here](#).
- `input_datasets=["images123"]` in the code above means that the script will consume all data in folder `images123` in blob storage as the input. The folder must exist in blob storage, in the location that you gave when creating the datastore. Once the script has run, it will also register the data in this folder as an AML dataset.
- `output_datasets=["images123_resized"]` means that the script will create a temporary folder when running in AML, and while the job writes data to that folder, upload it to blob storage, in the data store.

For more examples, please see [examples.md](#). For more details about datasets, see [here](#).

## 1.4 Additional arguments you should know about

`submit_to_azure_if_needed` has a large number of arguments, please check the API documentation for an exhaustive list. The particularly helpful ones are listed below.

- `experiment_name`: All runs in AzureML are grouped in “experiments”. By default, the experiment name is determined by the name of the script you submit, but you can specify a name explicitly with this argument.
- `environment_variables`: A dictionary with the contents of all environment variables that should be set inside the AzureML run, before the script is started.
- `docker_base_image`: This specifies the name of the Docker base image to use for creating the Python environment for your script. The amount of memory to allocate for Docker is given by `docker_shm_size`.
- `num_nodes`: The number of nodes on which your script should run. This is essential for distributed training.
- `tags`: A dictionary mapping from string to string, with additional tags that will be stored on the AzureML run. This is helpful to add metadata about the run for later use.

## 1.5 Conda environments, Alternate pips, Private wheels

The function `submit_to_azure_if_needed` tries to locate a Conda environment file in the current folder, or in the Python path, with the name `environment.yml`. The actual Conda environment file to use can be specified directly with:

```
run_info = submit_to_azure_if_needed(
    ...
    conda_environment_file=conda_environment_file,
```

where `conda_environment_file` is a `pathlib.Path` or a string identifying the Conda environment file to use.

The basic use of Conda assumes that packages listed are published [Conda packages](#) or published Python packages on [PyPI](#). However, during development, the Python package may be on [Test.PyPI](#), or in some other location, in which case the alternative package location can be specified directly with:

```
run_info = submit_to_azure_if_needed(
    ...
    pip_extra_index_url="https://test.pypi.org/simple/",
```

Finally, it is possible to use a private wheel, if the package is only available locally with:

```
run_info = submit_to_azure_if_needed(  
    ...  
    private_pip_wheel_path=private_pip_wheel_path,
```

where `private_pip_wheel_path` is a `pathlib.Path` or a string identifying the wheel package to use. In this case, this wheel will be copied to the AzureML environment as a private wheel.

## CONNECTING TO AZURE

### 2.1 Authentication

The `hi-ml` package uses two possible ways of authentication with Azure. The default is what is called “Interactive Authentication”. When you submit a job to Azure via `hi-ml`, this will use the credentials you used in the browser when last logging into Azure. If there are no credentials yet, you should see instructions printed out to the console about how to log in using your browser.

We recommend using Interactive Authentication.

Alternatively, you can use a so-called Service Principal, for example within build pipelines.

### 2.2 Service Principal Authentication

A Service Principal is a form of generic identity or machine account. This is essential if you would like to submit training runs from code, for example from within an Azure pipeline. You can find more information about application registrations and service principal objects [here](#).

If you would like to use Service Principal, you will need to create it in Azure first, and then store 3 pieces of information in 3 environment variables — please see the instructions below. When all the 3 environment variables are in place, your Azure submissions will automatically use the Service Principal to authenticate.

#### 2.2.1 Creating the Service Principal

1. Navigate back to [aka.ms/portal](https://aka.ms/portal)
2. Navigate to **App registrations** (use the top search bar to find it).
3. Click on + **New registration** on the top left of the page.
4. Choose a name for your application e.g. `MyServicePrincipal` and click **Register**.
5. Once it is created you will see your application in the list appearing under **App registrations**. This step might take a few minutes.
6. Click on the resource to access its properties. In particular, you will need the application ID. You can find this ID in the **Overview** tab (accessible from the list on the left of the page).
7. Create an environment variable called `HIML_SERVICE_PRINCIPAL_ID`, and set its value to the application ID you just saw.
8. You need to create an application secret to access the resources managed by this service principal. On the pane on the left find **Certificates & Secrets**. Click on + **New client secret** (bottom of the page), note down

your token. Warning: this token will only appear once at the creation of the token, you will not be able to re-display it again later.

9. Create an environment variable called `HIML_SERVICE_PRINCIPAL_PASSWORD`, and set its value to the token you just added.

## **2.2.2 Providing permissions to the Service Principal**

Now that your service principal is created, you need to give permission for it to access and manage your AzureML workspace. To do so:

1. Go to your AzureML workspace. To find it you can type the name of your workspace in the search bar above.
2. On the Overview page, there is a link to the Resource Group that contains the workspace. Click on that.
3. When on the Resource Group, navigate to **Access control**. Then click on **+ Add > Add role assignment**. A pane will appear on the the right. Select **Role > Contributor**. In the **Select** field type the name of your Service Principal and select it. Finish by clicking **Save** at the bottom of the pane.

## **2.2.3 Azure Tenant ID**

The last remaining piece is the Azure tenant ID, which also needs to be available in an environment variable. To get that ID:

1. Log into Azure
2. Via the search bar, find “Azure Active Directory” and open it.
3. In the overview of that, you will see a field “Tenant ID”
4. Create an environment variable called `HIML_TENANT_ID`, and set that to the tenant ID you just saw.

## DATASETS

### 3.1 Key concepts

We'll first outline a few concepts that are helpful for understanding datasets.

#### 3.1.1 Blob Storage

Firstly, there is [Azure Blob Storage](#). Each blob storage account has multiple containers - you can think of containers as big disks that store files. The `hi-ml` package assumes that your datasets live in one of those containers, and each top level folder corresponds to one dataset.

#### 3.1.2 AzureML Data Stores

Secondly, there are data stores. This is a concept coming from Azure Machine Learning, described [here](#). Data stores provide access to one blob storage account. They exist so that the credentials to access blob storage do not have to be passed around in the code - rather, the credentials are stored in the data store once and for all.

You can view all data stores in your AzureML workspace by clicking on one of the bottom icons in the left-hand navigation bar of the AzureML studio.

One of these data stores is designated as the default data store.

#### 3.1.3 AzureML Datasets

Thirdly, there are datasets. Again, this is a concept coming from Azure Machine Learning. A dataset is defined by

- A data store
- A set of files accessed through that data store

You can view all datasets in your AzureML workspace by clicking on one of the icons in the left-hand navigation bar of the AzureML studio.

### 3.1.4 Preparing data

To simplify usage, the `hi-ml` package creates AzureML datasets for you. All you need to do is to

- Create a blob storage account for your data, and within it, a container for your data.
- Create a data store that points to that storage account, and store the credentials for the blob storage account in it

From that point on, you can drop a folder of files in the container that holds your data. Within the `hi-ml` package, just reference the name of the folder, and the package will create a dataset for you, if it does not yet exist.

## 3.2 Using the datasets

The simplest way of specifying that your script uses a folder of data from blob storage is as follows: Add the `input_datasets` argument to your call of `submit_to_azure_if_needed` like this:

```
from health_azure import submit_to_azure_if_needed
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=["my_folder"],
                                     default_datastore="my_datastore")
input_folder = run_info.input_datasets[0]
```

What will happen under the hood?

- The toolbox will check if there is already an AzureML dataset called “my\_folder”. If so, it will use that. If there is no dataset of that name, it will create one from all the files in blob storage in folder “my\_folder”. The dataset will be created using the data store provided, “my\_datastore”.
- Once the script runs in AzureML, it will download the dataset “my\_folder” to a temporary folder.
- You can access this temporary location by `run_info.input_datasets[0]`, and read the files from it.

More complicated setups are described below.

### 3.2.1 Input and output datasets

Any run in AzureML can consume a number of input datasets. In addition, an AzureML run can also produce an output dataset (or even more than one).

Output datasets are helpful if you would like to run, for example, a script that transforms one dataset into another.

You can use that via the `output_datasets` argument:

```
from health_azure import submit_to_azure_if_needed
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=["my_folder"],
                                     output_datasets=["new_dataset"],
                                     default_datastore="my_datastore")
input_folder = run_info.input_datasets[0]
output_folder = run_info.output_datasets[0]
```

Your script can now read files from `input_folder`, transform them, and write them to `output_folder`. The latter will be a folder on the temp file system of the machine. At the end of the script, the contents of that temp folder will be uploaded to blob storage, and registered as a dataset.

### 3.2.2 Mounting and downloading

An input dataset can be downloaded before the start of the actual script run, or it can be mounted. When mounted, the files are accessed via the network once needed - this is very helpful for large datasets where downloads would create a long waiting time before the job start.

Similarly, an output dataset can be uploaded at the end of the script, or it can be mounted. Mounting here means that all files will be written to blob storage already while the script runs (rather than at the end).

Note: If you are using mounted output datasets, you should NOT rename files in the output folder.

Mounting and downloading can be triggered by passing in DatasetConfig objects for the input\_datasets argument, like this:

```
from health_azure import DatasetConfig, submit_to_azure_if_needed
input_dataset = DatasetConfig(name="my_folder", datastore="my_datastore", use_
    ↳mounting=True)
output_dataset = DatasetConfig(name="new_dataset", datastore="my_datastore", use_
    ↳mounting=True)
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=[input_dataset],
                                     output_datasets=[output_dataset])
input_folder = run_info.input_datasets[0]
output_folder = run_info.output_datasets[0]
```

### 3.2.3 Local execution

For debugging, it is essential to have the ability to run a script on a local machine, outside of AzureML. Clearly, your script needs to be able to access data in those runs too.

There are two ways of achieving that: Firstly, you can specify an equivalent local folder in the DatasetConfig objects:

```
from pathlib import Path
from health_azure import DatasetConfig, submit_to_azure_if_needed
input_dataset = DatasetConfig(name="my_folder",
                              datastore="my_datastore",
                              local_folder=Path("/datasets/my_folder_local"))
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=[input_dataset])
input_folder = run_info.input_datasets[0]
```

Secondly, you can check the returned path in run\_info, and replace it with something for local execution. run\_info.input\_datasets[0] will be None if the script runs outside of AzureML, and no local\_folder is available.

```
from pathlib import Path
from health_azure import submit_to_azure_if_needed
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=["my_folder"],
                                     default_datastore="my_datastore")
input_folder = run_info.input_datasets[0] or Path("/datasets/my_folder_local")
```

### 3.2.4 Making a dataset available at a fixed folder location

Occasionally, scripts expect the input dataset at a fixed location, for example, data is always read from `/tmp/mnist`. AzureML has the capability to download/mount a dataset to such a fixed location. With the `hi-ml` package, you can trigger that behaviour via an additional option in the `DatasetConfig` objects:

```
from health_azure import DatasetConfig, submit_to_azure_if_needed
input_dataset = DatasetConfig(name="my_folder",
                              datastore="my_datastore",
                              use_mounting=True,
                              target_folder="/tmp/mnist")
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=[input_dataset])
# Input_folder will now be "/tmp/mnist"
input_folder = run_info.input_datasets[0]
```

### 3.2.5 Dataset versions

AzureML datasets can have versions, starting at 1. You can view the different versions of a dataset in the AzureML workspace. In the `hi-ml` toolbox, you would always use the latest version of a dataset unless specified otherwise. If you do need a specific version, use the `version` argument in the `DatasetConfig` objects:

```
from health_azure import DatasetConfig, submit_to_azure_if_needed
input_dataset = DatasetConfig(name="my_folder",
                              datastore="my_datastore",
                              version=7)
run_info = submit_to_azure_if_needed(...,
                                     input_datasets=[input_dataset])
input_folder = run_info.input_datasets[0]
```



## HYPERPARAMETER SEARCH VIA HYPERDRIVE

`HyperDrive` runs can start multiple AzureML jobs in parallel. This can be used for tuning hyperparameters, or executing multiple training runs for cross validation. To use that with the `hi-ml` package, simply supply a `HyperDrive` configuration object as an additional argument. Note that this object needs to be created with an empty `run_config` argument (this will later be replaced with the correct `run_config` that submits your script.)

The example below shows a hyperparameter search that aims to minimize the validation loss `val_loss`, by choosing one of three possible values for the learning rate commandline argument `learning_rate`.

```
from azureml.core import ScriptRunConfig
from azureml.train.hyperdrive import GridParameterSampling, HyperDriveConfig, PrimaryMetricGoal, choice
from health_azure import submit_to_azure_if_needed
hyperdrive_config = HyperDriveConfig(
    run_config=ScriptRunConfig(source_directory=""),
    hyperparameter_sampling=GridParameterSampling(
        parameter_space={
            "learning_rate": choice([0.1, 0.01, 0.001])
        },
    primary_metric_name="val_loss",
    primary_metric_goal=PrimaryMetricGoal.MINIMIZE,
    max_total_runs=5
)
submit_to_azure_if_needed(..., hyperdrive_config=hyperdrive_config)
```

For further examples, please check the [example scripts here](#), and the [HyperDrive documentation](#).



## USING CHEAP LOW PRIORITY VMS

By using Low Priority machines in AzureML, we can run training at greatly reduced costs (around 20% of the original price, see references below for details). This comes with the risk, though, of having the job interrupted and later re-started. This document describes the inner workings of Low Priority compute, and how to best make use of it.

Because the jobs can get interrupted, low priority machines are not suitable for production workload where time is critical. They do offer a lot of benefits though for long-running training jobs or large scale experimentation, that would otherwise be expensive to carry out.

### 5.1 Setting up the Compute Cluster

Jobs in Azure Machine Learning run in a “compute cluster”. When creating a compute cluster, we can specify the size of the VM, the type and number of GPUs, etc. Doing this via the AzureML UI is described [here](#). Doing it programmatically is described [here](#).

One of the settings to tweak when creating the compute cluster is whether the machines are “Dedicated” or “Low Priority”:

- Dedicated machines will be permanently allocated to your compute cluster. The VMs in a dedicated cluster will be always available, unless the cluster is set up in a way that it removes idle machines. Jobs will not be interrupted.
- Low priority machines effectively make use of spare capacity in the data centers, you can think of them as “dedicated machines that are presently idle”. They are available at a much lower price (around 20% of the price of a dedicated machine). These machines are made available to you until they are needed as dedicated machines somewhere else.

In order to get a compute cluster that operates at the lowest price point, choose

- Low priority machines
- Set “Minimum number of nodes” to 0, so that the cluster removes all idle machines if no jobs are running.


For details on pricing, check [this Azure price calculator](#), choose “Category: GPU”. The price for low priority VMs is given in the “Spot” column

## 5.2 Behaviour of Low Priority VMs

Jobs can be interrupted at any point, this is called “low priority preemption”. When interrupted, the job stops - there is no signal that we can make use of to do cleanup or something. All the files that the job has produced up to that point in the `outputs` and `logs` folders will be saved to the cloud.

At some later point, the job will be assigned a virtual machine again. When re-started, all the files that the job had produced in its previous run will be available on disk again where they were before interruption, mounted at the same path. That is, if the interrupted job wrote a file `outputs/foo.txt`, this file will be accessible as `outputs/foo.txt` also after the restart.

Note that all AzureML-internal log files that the job produced in a previous run will be **overwritten** (this behaviour may change in the future). That is in contrast to the behaviour for metrics that the interrupted job had saved to AzureML already (for example, metrics written by a call like `Run.log("loss", loss_tensor.item())`): Those metrics are already stored in AzureML, and will still be there when the job restarts. The re-started job will then **append** to the metrics that had been written in the previous run. This typically shows as sudden jumps in metrics, as illustrated here:



– In this example, the learning rate was increasing for the first 6 or so epochs. Then the job got preempted, and started training from scratch, with the initial learning rate and schedule. Note that this behaviour is only an artifact of how the metrics are stored in AzureML, the actual training is doing the right thing.

How do you verify that your job got interrupted? Usually, you would see a warning displayed on the job page in the AzureML UI, that says something along the lines of “Low priority compute preemption warning: a node has been preempted.”. You can use kinks in metrics as another indicator that your job got preempted: Sudden jumps in metrics after which the metric follows a shape similar to the one at job start usually indicates low priority preemption.

Note that a job can be interrupted more than one time.

## 5.3 Best Practice Guide for Your Jobs

In order to make best use of low priority compute, your code needs to be made resilient to restarts. Essentially, this means that it should write regular checkpoints, and try to use those checkpoint files if they already exist. Examples of how to best do that are given below.

In addition, you need to bear in mind that the job can be interrupted at any moment, for example when it is busy uploading huge checkpoint files to Azure. When trying to upload again after restart, there can be resource collisions.

### 5.3.1 Writing and Using Recovery Checkpoints

When using PyTorch Lightning, you can add a checkpoint callback to your trainer, that ensures that you save the model and optimizer to disk in regular intervals. This callback needs to be added to your `Trainer` object. Note that these recovery checkpoints need to be written to the `outputs` folder, because only files in this folder get saved to Azure automatically when the job gets interrupted.

When starting training, your code needs to check if there is already a recovery checkpoint present on disk. If so, training should resume from that point.

Here is a code snippet that illustrates all that:

```
import re
from pathlib import Path
import numpy as np
from health_ml.utils import AzureMLLogger
from pytorch_lightning import Trainer
from pytorch_lightning.callbacks import ModelCheckpoint

RECOVERY_CHECKPOINT_FILE_NAME = "recovery_"
CHECKPOINT_FOLDER = "outputs/checkpoints"

def get_latest_recovery_checkpoint():
    all_recovery_files = [f for f in Path(CHECKPOINT_FOLDER).glob(RECOVERY_CHECKPOINT_
    FILE_NAME + "*")]
    if len(all_recovery_files) == 0:
        return None
    # Get recovery checkpoint with highest epoch number
    recovery_epochs = [int(re.findall(r"[\\d]+", f.stem)[0]) for f in all_recovery_files]
    idx_max_epoch = int(np.argmax(recovery_epochs))
    return str(all_recovery_files[idx_max_epoch])

recovery_checkpoint = ModelCheckpoint(dirpath=CHECKPOINT_FOLDER,
                                     filename=RECOVERY_CHECKPOINT_FILE_NAME + "{epoch}",
                                     period=10)
trainer = Trainer(default_root_dir="outputs",
                  callbacks=[recovery_checkpoint],
                  logger=[AzureMLLogger()],
                  resume_from_checkpoint=get_latest_recovery_checkpoint())
```

## 5.4 Additional Optimizers and Other State

In order to be resilient to interruption, your jobs need to save all their state to disk. In PyTorch Lightning training, this would include all optimizers that you are using. The “normal” optimizer for model training is saved to the checkpoint by Lightning already. However, you may be using callbacks or other components that maintain state. As an example, training a linear head for self-supervised learning can be done in a callback, and that callback can have its own optimizer. Such callbacks need to correctly implement the `on_save_checkpoint` method to save their state to the checkpoint, and `on_load_checkpoint` to load it back in.

For more information about persisting state, check the [PyTorch Lightning documentation](#) .

## COMMANDLINE TOOLS

### 6.1 Run TensorBoard

From the command line, run the command

```
himl-tb
```

specifying one of [--experiment] [--latest\_run\_file] [--run\_recovery\_ids] [--run\_ids]

This will start a TensorBoard session, by default running on port 6006. To use an alternative port, specify this with --port.

If --experiment is provided, the most recent Run from this experiment will be visualised. If --latest\_run\_file is provided, the script will expect to find a RunId in this file. Alternatively you can specify the Runs to visualise via --run\_recovery\_ids or --run\_ids.

By default, this tool expects that your TensorBoard logs live in a folder named 'logs' and will create a similarly named folder in your root directory. If your TensorBoard logs are stored elsewhere, you can specify this with the --log\_dir argument.

If you choose to specify --experiment, you can also specify --num\_runs to view and/or --tags to filter by.

If your AML config path is not ROOT\_DIR/config.json, you must also specify --config\_file.

To see an example of how to create TensorBoard logs using PyTorch on AML, see the AML submitting script which submits the following pytorch sample script. Note that to run this, you'll need to create an environment with pytorch and tensorboard as dependencies, as a minimum. See an example conda environemnt. This will create an experiment named 'tensorboard\_test' on your Workspace, with a single run. Go to outputs + logs -> outputs to see the tensorboard events file.

### 6.2 Download files from AML Runs

From the command line, run the command

```
himl-download
```

specifying one of [--experiment] [--latest\_run\_file] [--run\_recovery\_ids] [--run\_ids]

If --experiment is provided, the most recent Run from this experiment will be downloaded. If --latest\_run\_file is provided, the script will expect to find a RunId in this file. Alternatively you can specify the Run to download via --run\_recovery\_ids or --run\_ids.

The files associated with your Run will be downloaded to the location specified with --output\_dir (by default ROOT\_DIR/outputs)

If you choose to specify --experiment, you can also specify --tags to filter by.

If your AML config path is not `ROOT_DIR/config.json`, you must also specify `--config_file`.



## DOWNLOADING FROM/ UPLOADING TO AZURE ML

All of the below functions will attempt to find a current workspace, if running in Azure ML, or else will attempt to locate 'config.json' file in the current directory, and its parents. Alternatively, you can specify your own Workspace object or a path to a file containing the workspace settings.

### 7.1 Download files from an Azure ML Run

To download all files from an AML Run, given its run id, perform the following:

```
from pathlib import Path
from health_azure import download_files_from_run_id
run_id = "example_run_id_123"
output_folder = Path("path/to/save")
download_files_from_run_id(run_id, output_folder)
```

Here, "path\_to\_save" represents the folder in which we want the downloaded files to be stored. E.g. if your run contains the files ["abc/def/1.txt", "abc/2.txt"] and you specify the prefix "abc" and the output\_folder "my\_outputs", you'll end up with the files ["my\_outputs/abc/def/1.txt", "my\_outputs/abc/2.txt"]

If you wish to specify the file name(s) to be downloaded, you can do so with the "prefix" parameter. E.g. prefix="outputs" will download all files within the "output" folder, if such a folder exists within your Run.

There is an additional parameter, "validate\_checksum" which defaults to False. If True, will validate MD5 hash of the data arriving (in chunks) to that being sent.

Note that if your code is running in a distributed manner, files will only be downloaded onto nodes with local rank = 0. E.g. if you have 2 nodes each running 4 processes, the file will be downloaded by CPU/GPU 0 on each of the 2 nodes. All processes will be synchronized to only exit the downloading method once it has completed on all nodes/ranks.

### 7.2 Downloading checkpoint files from a run

To download checkpoint files from an Azure ML Run, perform the following:

```
from pathlib import Path
from health_azure import download_checkpoints_from_run_id
download_checkpoints_from_run_id("example_run_id_123", Path("path/to/checkpoint/directory
↪"))
```

All files within the checkpoint directory will be downloaded into the folder specified by "path/to/checkpoint\_directory".

Since checkpoint files are often large and therefore prone to corruption during download, by default, this function will validate the MD5 hash of the data downloaded (in chunks) compared to that being sent.

Note that if your code is running in a distributed manner, files will only be downloaded onto nodes with local rank = 0. E.g. if you have 2 nodes each running 4 processes, the file will be downloaded by CPU/GPU 0 on each of the 2 nodes. All processes will be synchronized to only exit the downloading method once it has completed on all nodes/ranks.

## 7.3 Downloading files from an Azure ML Datastore

To download data from an Azure ML Datastore within your Workspace, follow this example:

```
from pathlib import Path
from health_azure import download_from_datastore
download_from_datastore("datastore_name", "prefix", Path("path/to/output/directory") )
```

where “prefix” represents the path to the file(s) to be downloaded, relative to the datastore “datastore\_name”. Azure will search for files within the Datastore whose paths begin with this string. If you wish to download multiple files from the same folder, set equal to that folder’s path within the Datastore. If you wish to download a single file, include both the path to the folder it resides in, as well as the filename itself. If the relevant file(s) are found, they will be downloaded to the folder specified by <output\_folder>. If this directory does not already exist, it will be created. E.g. if your datastore contains the paths [“foo/bar/1.txt”, “foo/bar/2.txt”] and you call this function with file\_prefix=“foo/bar” and output\_folder=“outputs”, you would end up with the files [“outputs/foo/bar/1.txt”, “outputs/foo/bar/2.txt”]

This function takes additional parameters “overwrite” and “show\_progress”. If True, overwrite will overwrite any existing local files with the same path. If False and there is a duplicate file, it will skip this file. If show\_progress is set to True, the progress of the file download will be visible in the terminal.

## 7.4 Uploading files to an Azure ML Datastore

To upload data to an Azure ML Datastore within your workspace, perform the following:

```
from pathlib import Path
from health_azure import upload_to_datastore
upload_to_datastore("datastore_name", Path("path/to/local/data/folder"), Path("path/to/
↳datastore/folder") )
```

Where “datastore\_name” is the name of the registered Datastore within your workspace that you wish to upload to and “path/to/datastore/folder” is the relative path within this Datastore that you wish to upload data to. Note that the path to local data must be a folder, not a single path. The folder name will not be included in the remote path. E.g. if you specify the local\_data\_dir=“foo/bar” and that contains the files [“1.txt”, “2.txt”], and you specify the remote\_path=“baz”, you would see the following paths uploaded to your Datastore: [“baz/1.txt”, “baz/2.txt”]

This function takes additional parameters “overwrite” and “show\_progress”. If True, overwrite will overwrite any existing remote files with the same path. If False and there is a duplicate file, it will skip this file. If show\_progress is set to True, the progress of the file upload will be visible in the terminal.

## EXAMPLES

**Note:** All examples below contain links to sample scripts that are also included in the repository. The experience is **optimized for use on readthedocs**. When navigating to the sample scripts on the github UI, you will only see the `.rst` file that links to the `.py` file. To access the `.py` file, go to the folder that contains the respective `.rst` file.

### 8.1 Basic integration

The sample `examples/1/sample.py` is a script that takes an optional command line argument of a target value and prints all the prime numbers up to (but not including) this target. It is simply intended to demonstrate a long running operation that we want to run in Azure. Run it using e.g.

```
cd examples/1
python sample.py -n 103
```

The sample `examples/2/sample.py` shows the minimal modifications to run this in AzureML. Firstly create an AzureML workspace and download the config file, as explained [here](#). The config file should be placed in the same folder as the sample script. A sample Conda environment file is supplied. Import the [hi-ml package](#) into the current environment. Finally add the following to the sample script:

```
from health_azure import submit_to_azure_if_needed
...
def main() -> None:
    _ = submit_to_azure_if_needed(
        compute_cluster_name="lite-testing-ds2",
        wait_for_completion=True,
        wait_for_completion_show_output=True)
```

Replace `lite-testing-ds2` with the name of a compute cluster created within the AzureML workspace. If this script is invoked as the first sample, e.g.

```
cd examples/2
python sample.py -n 103
```

then the output will be exactly the same. But if the script is invoked as follows:

```
cd examples/2
python sample.py -n 103 --azureml
```

then the function `submit_to_azure_if_needed` will perform all the required actions to run this script in AzureML and exit. Note that:

- code after `submit_to_azure_if_needed` is not run locally, but it is run in AzureML.

- the print statement prints to the AzureML console output and is available in the Output + logs tab of the experiment in the 70\_driver\_log.txt file, and can be downloaded from there.
- the command line arguments are passed through (apart from `-azureml`) when running in AzureML.
- a new file: `most_recent_run.txt` will be created containing an identifier of this AzureML run.

A sample script `examples/2/results.py` demonstrates how to programmatically download the driver log file.

## 8.2 Output files

The sample `examples/3/sample.py` demonstrates output file handling when running on AzureML. Because each run is performed in a separate VM or cluster then any file output is not generally preserved. In order to keep the output it should be written to the `outputs` folder when running in AzureML. The AzureML infrastructure will preserve this and it will be available for download from the `outputs` folder in the Output + logs tab.

Make the following additions:

```
from health_azure import submit_to_azure_if_needed
run_info = submit_to_azure_if_needed(
    ...
    parser.add_argument("-o", "--output", type=str, default="primes.txt", required=False,
    ↪ help="Output file name")
    ...
    output = run_info.output_folder / args.output
    output.write_text("\n".join(map(str, primes)))
```

When running locally `submit_to_azure_if_needed` will create a subfolder called `outputs` and then the output can be written to the file `args.output` there. When running in AzureML the output will be available in the file `args.output` in the Experiment.

A sample script `examples/3/results.py` demonstrates how to programmatically download the output file.

## 8.3 Output datasets

The sample `examples/4/sample.py` demonstrates output dataset handling when running on AzureML.

In this case, the following parameters are added to `submit_to_azure_if_needed`:

```
from health_azure import submit_to_azure_if_needed
run_info = submit_to_azure_if_needed(
    ...
    default_datastore="himldatasets",
    output_datasets=["himl_sample4_output"],
```

The `default_datastore` is required if using the simplest configuration for an output dataset, to just use the blob container name. There is an alternative that doesn't require the `default_datastore` and allows a different datastore for each dataset:

```
from health_azure import DatasetConfig, submit_to_azure_if_needed
...
run_info = submit_to_azure_if_needed(
    ...
    output_datasets=[DatasetConfig(name="himl_sample4_output", datastore=
    ↪ "himldatasets")]
```

(continues on next page)

(continued from previous page)

Now the output folder is constructed as follows:

```
output_folder = run_info.output_datasets[0] or Path("outputs") / "himl_sample4_output"
↪
output_folder.mkdir(parents=True, exist_ok=True)
output = output_folder / args.output
```

When running in AzureML `run_info.output_datasets[0]` will be populated using the new parameter and the output will be written to that blob storage. When running locally `run_info.output_datasets[0]` will be `None` and a local folder will be created and used.

A sample script `examples/4/results.py` demonstrates how to programmatically download the output dataset file.

For more details about datasets, see [here](#)

## 8.4 Input datasets

This example trains a simple classifier on a toy dataset, first creating the dataset files and then in a second script training the classifier.

The script `examples/5/inputs.py` is provided to prepare the csv files. Run the script to download the Iris dataset and create two CSV files:

```
cd examples/5
python inputs.py
```

The training script `examples/5/sample.py` is modified from [https://github.com/Azure/MachineLearningNotebooks/blob/master/how-to-use-azureml/ml-frameworks/scikit-learn/train-hyperparameter-tune-deploy-with-sklearn/train\\_iris.py](https://github.com/Azure/MachineLearningNotebooks/blob/master/how-to-use-azureml/ml-frameworks/scikit-learn/train-hyperparameter-tune-deploy-with-sklearn/train_iris.py) to work with input csv files. Start it to train the actual classifier, based on the data files that were just written:

```
cd examples/5
python sample.py
```

### 8.4.1 Including input files in the snapshot

When using very small datafiles (in the order of few MB), the easiest way to get the input data to Azure is to include them in the set of (source) files that are uploaded to Azure. You can run the dataset creation script on your local machine, writing the resulting two files to the same folder where your training script is located, and then submit the training script to AzureML. Because the dataset files are in the same folder, they will automatically be uploaded to AzureML.

However, it is not ideal to have the input files in the snapshot: The size of the snapshot is limited to 25 MB. It is better to put the data files into blob storage and use input datasets.

## 8.4.2 Creating the dataset in AzureML

The suggested way of creating a dataset is to run a script in AzureML that writes an output dataset. This is particularly important for large datasets, to avoid the usually low bandwidth from a local machine to the cloud.

This is shown in `examples/6/inputs.py`: This script prepares the CSV files in an AzureML run, and writes them to an output dataset called `himl_sample6_input`. The relevant code parts are:

```
run_info = submit_to_azure_if_needed(
    compute_cluster_name="lite-testing-ds2",
    default_datastore="himldatastore",
    output_datasets=["himl_sample6_input"])
# The dataset files should be written into this folder:
dataset = run_info.output_datasets[0] or Path("dataset")
```

Run the script:

```
cd examples/6
python inputs.py --azureml
```

You can now modify the training script `examples/6/sample.py` to use the newly created dataset `himl_sample6_input` as an input. To do that, the following parameters are added to `submit_to_azure_if_needed`:

```
run_info = submit_to_azure_if_needed(
    compute_cluster_name="lite-testing-ds2",
    default_datastore="himldatastore",
    input_datasets=["himl_sample6_input"])
```

When running in AzureML, the dataset will be downloaded before running the job. You can access the temporary folder where the dataset is available like this:

```
input_folder = run_info.input_datasets[0] or Path("dataset")
```

The part behind the `or` statement is only necessary to keep a reasonable behaviour when running outside of AzureML: When running in AzureML `run_info.input_datasets[0]` will be populated using input dataset specified in the call to `submit_to_azure_if_needed`, and the input will be downloaded from blob storage. When running locally `run_info.input_datasets[0]` will be `None` and a local folder should be populated and used.

The `default_datastore` is required if using the simplest configuration for an input dataset. There are alternatives that do not require the `default_datastore` and allows a different datastore for each dataset, for example:

```
from health_azure import DatasetConfig, submit_to_azure_if_needed
...
run_info = submit_to_azure_if_needed(
    ...
    input_datasets=[DatasetConfig(name="himl_sample7_input", datastore="himldatastore"
    ↪)],
```

For more details about datasets, see [here](#)

### 8.4.3 Uploading the input files manually

An alternative to writing the dataset in AzureML (as suggested above) is to create them on the local machine, and upload them manually directly to Azure blob storage.

This is shown in `examples/7/inputs.py`: This script prepares the CSV files and uploads them to blob storage, in a folder called `himl_sample7_input`. Run the script:

```
cd examples/7
python inputs_via_upload.py
```

As in the above example, you can now modify the training script `examples/7/sample.py` to use an input dataset that has the same name as the folder where the files just got uploaded. In this case, the following parameters are added to `submit_to_azure_if_needed`:

```
run_info = submit_to_azure_if_needed(
    ...
    default_datastore="himldatasets",
    input_datasets=["himl_sample7_input"],
```

## 8.5 Hyperdrive

The sample `examples/8/sample.py` demonstrates adding hyperparameter tuning. This shows the same hyperparameter search as in the [AzureML sample](#).

Make the following additions:

```
from azureml.core import ScriptRunConfig
from azureml.train.hyperdrive import HyperDriveConfig, PrimaryMetricGoal, choice
from azureml.train.hyperdrive.sampling import RandomParameterSampling

...
def main() -> None:
    param_sampling = RandomParameterSampling({
        "--kernel": choice('linear', 'rbf', 'poly', 'sigmoid'),
        "--penalty": choice(0.5, 1, 1.5)
    })

    hyperdrive_config = HyperDriveConfig(
        run_config=ScriptRunConfig(source_directory=""),
        hyperparameter_sampling=param_sampling,
        primary_metric_name='Accuracy',
        primary_metric_goal=PrimaryMetricGoal.MAXIMIZE,
        max_total_runs=12,
        max_concurrent_runs=4)

    run_info = submit_to_azure_if_needed(
        ...
        hyperdrive_config=hyperdrive_config)
```

Note that this does not make sense to run locally, it should always be run in AzureML. When invoked with:

```
cd examples/8
python sample.py --azureml
```

this will perform a Hyperdrive run in AzureML, i.e. there will be 12 child runs, each randomly drawing from the parameter sample space. AzureML can plot the metrics from the child runs, but to do that, some small modifications are required.

Add in:

```
run = run_info.run
...
args = parser.parse_args()
run.log('Kernel type', np.str(args.kernel))
run.log('Penalty', np.float(args.penalty))
...
print('Accuracy of SVM classifier on test set: {:.2f}'.format(accuracy))
run.log('Accuracy', np.float(accuracy))
```

and these metrics will be displayed on the child runs tab in the Experiment page on AzureML.

## 8.6 Controlling when to submit to AzureML and when not

By default, the `hi-ml` package assumes that you supply a commandline argument `--azureml` (that can be anywhere on the commandline) to trigger a submission of the present script to AzureML. If you wish to control it via a different flag, coming out of your own argument parser, use the `submit_to_azureml` argument of the function `health.azure.himl.submit_to_azure_if_needed`.



## LOGGING METRICS WHEN TRAINING MODELS IN AZUREML

This section describes the basics of logging to AzureML, and how this can be simplified when using PyTorch Lightning. It also describes helper functions to make logging more consistent across your code.

### 9.1 Basics

The mechanics of writing metrics to an ML training run inside of AzureML are described [here](#).

Using the `hi-ml-azure` toolbox, you can simplify that like this:

```
from health_azure import RUN_CONTEXT
...
RUN_CONTEXT.log(name="name_of_the_metric", value=my_tensor.item())
```

Similarly you can log strings (via the `log_text` method) or figures (via the `log_image` method), see the [documentation](#).

### 9.2 Using PyTorch Lightning

The `hi-ml` toolbox relies on `pytorch-lightning` for a lot of its functionality. Logging of metrics is described in detail [here](#)

`hi-ml` provides a Lightning-ready logger object to use with AzureML. You can add that to your trainer as you would add a Tensorboard logger, and afterwards see all metrics in both your Tensorboard files and in the AzureML UI. This logger can be added to the Trainer object as follows:

```
from health_ml.utils import AzureMLLogger
from pytorch_lightning.loggers import TensorBoardLogger
tb_logger = TensorBoardLogger("logs/")
azureml_logger = AzureMLLogger()
trainer = Trainer(logger=[tb_logger, azureml_logger])
```

You do not need to make any changes to your logging code to write to both loggers at the same time. This means that, if your code correctly writes to Tensorboard in a local run, you can expect the metrics to come out correctly in the AzureML UI as well after adding the `AzureMLLogger`.

## 9.3 Making logging consistent when training with PyTorch Lightning

A common problem of training scripts is that the calls to the logging methods tend to run out of sync. The `.log` method of a `LightningModule` has a lot of arguments, some of which need to be set correctly when running on multiple GPUs.

To simplify that, there is a function `log_on_epoch` that turns synchronization across nodes on/off depending on the number of GPUs, and always forces the metrics to be logged upon epoch completion. Use as follows:

```
from health_ml.utils import log_on_epoch
from pytorch_lightning import LightningModule

class MyModule(LightningModule):
    def training_step(self, *args, **kwargs):
        ...
        loss = my_loss(y_pred, y)
        log_on_epoch(self, loss)
        return loss
```

### 9.3.1 Logging learning rates

Logging learning rates is important for monitoring training, but again this can add overhead. To log learning rates easily and consistently, we suggest either of two options:

- Add a `LearningRateMonitor` callback to your trainer, as described [here](#)
- Use the `hi-ml` function `log_learning_rate`

The `log_learning_rate` function can be used at any point the training code, like this:

```
from health_ml.utils import log_learning_rate
from pytorch_lightning import LightningModule

class MyModule(LightningModule):
    def training_step(self, *args, **kwargs):
        ...
        log_learning_rate(self, "learning_rate")
        loss = my_loss(y_pred, y)
        return loss
```

`log_learning_rate` will log values from all learning rate schedulers, and all learning rates if a scheduler returns multiple values. In this example, the logged metric will be `learning_rate` if there is a single scheduler that outputs a single LR, or `learning_rate/1/0` to indicate the value coming from scheduler index 1, value index 0.

## NOTES FOR DEVELOPERS

### 10.1 Creating a Conda environment

To create a separate Conda environment with all packages that `hi-ml` requires for running and testing, use the provided `environment.yml` file. Create a Conda environment called `himl` from that via

```
conda env create --file environment.yml
conda activate himl
```

### 10.2 Using specific versions `hi-ml` in your Python environments

If you'd like to test specific changes to the `hi-ml` package in your code, you can use two different routes:

- You can clone the `hi-ml` repository on your machine, and use `hi-ml` in your Python environment via a local package install:

```
pip install -e <your_git_folder>/hi-ml
```

- You can consume an early version of the package from `test.pypi.org` via `pip`:

```
pip install --extra-index-url https://test.pypi.org/simple/ hi-ml==0.1.0.post165
```

- If you are using Conda, you can add an additional parameter for `pip` into the `Conda environment.yml` file like this:

```
name: foo
dependencies:
  - pip=20.1.1
  - python=3.7.3
  - pip:
    - --extra-index-url https://test.pypi.org/simple/
    - hi-ml==0.1.0.post165
```

## 10.3 Common things to do

The repository contains a makefile with definitions for common operations.

- `make check`: Run `flake8` and `mypy` on the repository.
- `make test`: Run `flake8` and `mypy` on the repository, then all tests via `pytest`
- `make pip`: Install all packages for running and testing in the current interpreter.
- `make conda`: Update the hi-ml Conda environment and activate it

## 10.4 Building documentation

To build the sphinx documentation, you must have sphinx and related packages installed (see `build_requirements.txt` in the repository root). Then run:

```
cd docs
make html
```

This will build all your documentation in `docs/build/html`.

## 10.5 Setting up your AzureML workspace

- In the browser, navigate to the AzureML workspace that you want to use for running your tests.
- In the top right section, there will be a dropdown menu showing the name of your AzureML workspace. Expand that.
- In the panel, there is a link “Download config file”. Click that.
- This will download a file `config.json`. Move that file to the root folder of your hi-ml repository. The file name is already present in `.gitignore`, and will hence not be checked in.

## 10.6 Creating and Deleting Docker Environments in AzureML

- Passing a `docker_base_image` into `submit_to_azure_if_needed` causes a new image to be built and registered in your workspace (see [docs](#) for more information).
- To remove an environment use the `az ml environment delete` function in the AzureML CLI (note that all the parameters need to be set, none are optional).

## 10.7 Testing

For all of the tests to work locally you will need to cache your AzureML credentials. One simple way to do this is to run the example in `src/health/azure/examples` (i.e. run `python elevate_this.py --message='Hello World' --azureml` or `make example`) after editing `elevate_this.py` to reference your compute cluster.

When running the tests locally, they can either be run against the source directly, or the source built into a package.

- To run the tests against the source directly in the local `src` folder, ensure that there is no wheel in the `dist` folder (for example by running `make clean`). If a wheel is not detected, then the local `src` folder will be copied into the temporary test folder as part of the test process.
- To run the tests against the source as a package, build it with `make build`. This will build the local `src` folder into a new wheel in the `dist` folder. This wheel will be detected and passed to AzureML as a private package as part of the test process.

## 10.8 Creating a New Release

To create a new package release, follow these steps:

- Modify `CHANGELOG.md` as follows:
  - Copy the whole section called “Upcoming”, including its subsections for “Added/Removed/...” to a new section, that has the desired package version and the current date as the title. For example, to release package version `0.12.17` on Oct 10th, that section would be called “`0.12.17 (2021-10-21)`”.
  - In the section for the new release, remove any empty subsections if needed.
  - Clean up all PR links from the “Upcoming” section, to effectively create an empty template for the next release.
- Create a PR for this change. While creating the PR, add the “no changelog needed” label that exists on the repo. *Important:* This label needs to be added right when the PR is created, not afterwards - the github workflows will not pick it up if added afterwards. In the worst case, you can add the label afterwards, and push a whitespace change to the PR.
- Once the PR with the updated `CHANGELOG.md` is in, create a tag that has the desired version number, plus a “v” prefix. For example, to create package version `0.12.17`, create a tag `v0.12.17`



## CONTRIBUTING TO THIS TOOLBOX

We welcome all contributions that help us achieve our aim of speeding up ML/AI research in health and life sciences. Examples of contributions are

- Data loaders for specific health & life sciences data
- Network architectures and components for deep learning models
- Tools to analyze and/or visualize data

All contributions to the toolbox need to come with unit tests, and will be reviewed when a Pull Request (PR) is started. If in doubt, reach out to the core `hi-ml` team before starting your work.

Please look through the existing folder structure to find a good home for your contribution.

### 11.1 Submitting a Pull Request

If you'd like to submit a PR to the codebase, please ensure you:

- Include a brief description
- Link to an issue, if relevant
- Write unit tests for the code - see below for details.
- Add appropriate documentation for any new code that you introduce

### 11.2 Code style

- We use `flake8` as a linter, and `mypy` for static typechecking. Both tools run as part of the PR build, and must run without errors for a contribution to be accepted. `mypy` requires that all functions and methods carry type annotations, see [mypy documentation](#).
- We highly recommend to run both tools *before* pushing the latest changes to a PR. If you have `make` installed, you can run both tools in one go via `make check` (from the repository root folder)
- Code should use sphinx-style comments like this:

```
from typing import List, Optional
def foo(bar: int) -> Optional[List]:
    """
    Creates a list. Or not. Note that there must a be blank line after the summary.

    :param bar: The length of the list. If 0, returns None. If there is a very long doc_
    ↪ string for an argument, the next
```

(continues on next page)

(continued from previous page)

```
    line must be indented.  
:return: A list with `bar` elements.  
"""
```

Note the blank line after the summary, and the indentation of multi-line documentation for parameters.

## 11.3 Unit testing

- DO write unit tests for each new function or class that you add.
- DO extend unit tests for existing functions or classes if you change their core behaviour.
- DO try your best to write unit tests that are fast. Very often, this can be done by reducing data size to a minimum. Also, it is helpful to avoid long-running integration tests, but try to test at the level of the smallest involved function.
- DO ensure that your tests are designed in a way that they can pass on the local machine, even if they are relying on specific cloud features. If required, use `unittest.mock` to simulate the cloud features, and hence enable the tests to run successfully on your local machine.
- DO run all unit tests on your dev machine before submitting your changes. The test suite is designed to pass completely also outside of cloud builds.
- DO NOT rely only on the test builds in the cloud (i.e., run test locally before submitting). Cloud builds trigger AzureML runs on GPU machines that have a far higher CO2 footprint than your dev machine.



## HEALTH\_AZURE PACKAGE

### 12.1 Functions

<i>create_run_configuration</i> (workspace, ...[, ...])	Creates an AzureML run configuration, that contains information about environment, multi node execution, and Docker.
<i>create_script_run</i> ([snapshot_root_directory, ...])	Creates an AzureML ScriptRunConfig object, that holds the information about the snapshot, the entry script, and its arguments.
<i>download_files_from_run_id</i> (run_id, output_folder)	For a given Azure ML run id, first retrieve the Run, and then download all files, which optionally start with a given prefix.
<i>download_checkpoints_from_run_id</i> (run_id, ...)	Given an Azure ML run id, download all files from a given checkpoint directory within that run, to the path specified by output_path.
<i>download_from_datastore</i> (datastore_name, ...)	Download file(s) from an Azure ML Datastore that are registered within a given Workspace.
<i>fetch_run</i> (workspace, run_recovery_id)	Finds an existing run in an experiment, based on a recovery ID that contains the experiment ID and the actual RunId.
<i>get_most_recent_run</i> (run_recovery_file, workspace)	Gets the name of the most recently executed AzureML run, instantiates that Run object and returns it.
<i>get_workspace</i> (aml_workspace, ...)	Retrieve an Azure ML Workspace from one of several places:
<i>is_running_in_azure_ml</i> ([aml_run])	Returns True if the given run is inside of an AzureML machine, or False if it is on a machine outside AzureML.
<i>set_environment_variables_for_multi_node</i> ()	Sets the environment variables that PyTorch Lightning needs for multi-node training.
<i>split_recovery_id</i> (id)	Splits a run ID into the experiment name and the actual run.
<i>submit_run</i> (workspace, experiment_name, ...)	Starts an AzureML run on a given workspace, via the script_run_config.
<i>submit_to_azure_if_needed</i> ([...])	Submit a folder to Azure, if needed and run it.
<i>torch_barrier</i> ()	This is a barrier to use in distributed jobs.
<i>upload_to_datastore</i> (datastore_name, ...[, ...])	Upload a folder to an Azure ML Datastore that is registered within a given Workspace.

### 12.1.1 create\_run\_configuration

```
health_azure.create_run_configuration(workspace, compute_cluster_name,
                                     conda_environment_file=None, aml_environment_name="",
                                     environment_variables=None, pip_extra_index_url="",
                                     private_pip_wheel_path=None, docker_base_image="",
                                     docker_shm_size="", num_nodes=1, max_run_duration="",
                                     input_datasets=None, output_datasets=None)
```

Creates an AzureML run configuration, that contains information about environment, multi node execution, and Docker.

#### Parameters

- **workspace** (Workspace) – The AzureML Workspace to use.
- **aml\_environment\_name** (str) – The name of an AzureML environment that should be used to submit the script. If not provided, an environment will be created from the arguments to this function (conda\_environment\_file, pip\_extra\_index\_url, environment\_variables, docker\_base\_image)
- **max\_run\_duration** (str) – The maximum runtime that is allowed for this job in AzureML. This is given as a floating point number with a string suffix s, m, h, d for seconds, minutes, hours, day. Examples: '3.5h', '2d'
- **compute\_cluster\_name** (str) – The name of the AzureML cluster that should run the job. This can be a cluster with CPU or GPU machines.
- **conda\_environment\_file** (Optional[Path]) – The conda configuration file that describes which packages are necessary for your script to run.
- **environment\_variables** (Optional[Dict[str, str]]) – The environment variables that should be set when running in AzureML.
- **docker\_base\_image** (str) – The Docker base image that should be used when creating a new Docker image.
- **docker\_shm\_size** (str) – The Docker shared memory size that should be used when creating a new Docker image.
- **pip\_extra\_index\_url** (str) – If provided, use this PIP package index to find additional packages when building the Docker image.
- **private\_pip\_wheel\_path** (Optional[Path]) – If provided, add this wheel as a private package to the AzureML workspace.
- **conda\_environment\_file** – The file that contains the Conda environment definition.
- **input\_datasets** (Optional[List[DatasetConfig]]) – The script will consume all data in folder in blob storage as the input. The folder must exist in blob storage, in the location that you gave when creating the datastore. Once the script has run, it will also register the data in this folder as an AzureML dataset.
- **output\_datasets** (Optional[List[DatasetConfig]]) – The script will create a temporary folder when running in AzureML, and while the job writes data to that folder, upload it to blob storage, in the data store.
- **num\_nodes** (int) – The number of nodes to use in distributed training on AzureML.

**Return type** RunConfiguration

**Returns**

### 12.1.2 create\_script\_run

`health_azure.create_script_run(snapshot_root_directory=None, entry_script=None, script_params=None)`  
Creates an AzureML ScriptRunConfig object, that holds the information about the snapshot, the entry script, and its arguments.

#### Parameters

- **entry\_script** (Union[Path, str, None]) – The script that should be run in AzureML.
- **snapshot\_root\_directory** (Optional[Path]) – The directory that contains all code that should be packaged and sent to AzureML. All Python code that the script uses must be copied over.
- **script\_params** (Optional[List[str]]) – A list of parameter to pass on to the script as it runs in AzureML. If empty (or None, the default) these will be copied over from sys.argv, omitting the `–azureml` flag.

**Return type** ScriptRunConfig

**Returns**

### 12.1.3 download\_files\_from\_run\_id

`health_azure.download_files_from_run_id(run_id, output_folder, prefix="", workspace=None, workspace_config_path=None, validate_checksum=False)`

For a given Azure ML run id, first retrieve the Run, and then download all files, which optionally start with a given prefix. E.g. if the Run creates a folder called “outputs”, which you wish to download all files from, specify `prefix="outputs"`. To download all files associated with the run, leave prefix empty.

If not running inside AML and neither a workspace nor the config file are provided, the code will try to locate a `config.json` file in any of the parent folders of the current working directory. If that succeeds, that `config.json` file will be used to instantiate the workspace.

If function is called in a distributed PyTorch training script, the files will only be downloaded once per node (i.e., all process where `is_local_rank_zero() == True`). All processes will exit this function once all downloads are completed.

#### Parameters

- **run\_id** (str) – The id of the Azure ML Run
- **output\_folder** (Path) – Local directory to which the Run files should be downloaded.
- **prefix** (str) – Optional prefix to filter Run files by
- **workspace** (Optional[Workspace]) – Optional Azure ML Workspace object
- **workspace\_config\_path** (Optional[Path]) – Optional path to settings for Azure ML Workspace
- **validate\_checksum** (bool) – Whether to validate the content from HTTP response

**Return type** None

### 12.1.4 download\_checkpoints\_from\_run\_id

`health_azure.download_checkpoints_from_run_id(run_id, checkpoint_dir, output_folder, aml_workspace=None, workspace_config_path=None)`

Given an Azure ML run id, download all files from a given checkpoint directory within that run, to the path specified by `output_path`. If running in AML, will take the current workspace. Otherwise, if neither `aml_workspace` nor `workspace_config_path` are provided, will try to locate a `config.json` file in any of the parent folders of the current working directory.

#### Parameters

- **run\_id** (str) – The id of the run to download checkpoints from
- **checkpoint\_dir** (str) – The path to the checkpoints directory within the run files
- **output\_folder** (Path) – The path to which the checkpoints should be stored
- **aml\_workspace** (Optional[Workspace]) – Optional AML workspace object
- **workspace\_config\_path** (Optional[Path]) – Optional workspace config file

**Return type** None

### 12.1.5 download\_from\_datastore

`health_azure.download_from_datastore(datastore_name, file_prefix, output_folder, aml_workspace=None, workspace_config_path=None, overwrite=False, show_progress=False)`

Download file(s) from an Azure ML Datastore that are registered within a given Workspace. The path to the file(s) to be downloaded, relative to the datastore `<datastore_name>`, is specified by the parameter “`prefix`”. Azure will search for files within the Datastore whose paths begin with this string. If you wish to download multiple files from the same folder, set `<prefix>` equal to that folder’s path within the Datastore. If you wish to download a single file, include both the path to the folder it resides in, as well as the filename itself. If the relevant file(s) are found, they will be downloaded to the folder specified by `<output_folder>`. If this directory does not already exist, it will be created. E.g. if your datastore contains the paths [“foo/bar/1.txt”, “foo/bar/2.txt”] and you call this function with `file_prefix=“foo/bar”` and `output_folder=“outputs”`, you would end up with the files [“outputs/foo/bar/1.txt”, “outputs/foo/bar/2.txt”]

If not running inside AML and neither a workspace nor the config file are provided, the code will try to locate a `config.json` file in any of the parent folders of the current working directory. If that succeeds, that `config.json` file will be used to instantiate the workspace.

#### Parameters

- **datastore\_name** (str) – The name of the Datastore containing the blob to be downloaded. This Datastore itself must be an instance of an `AzureBlobDatastore`.
- **file\_prefix** (str) – The prefix to the blob to be downloaded
- **output\_folder** (Path) – The directory into which the blob should be downloaded
- **aml\_workspace** (Optional[Workspace]) – Optional Azure ML Workspace object
- **workspace\_config\_path** (Optional[Path]) – Optional path to settings for Azure ML Workspace
- **overwrite** (bool) – If True, will overwrite any existing file at the same remote path. If False, will skip any duplicate file.
- **show\_progress** (bool) – If True, will show the progress of the file download

**Return type** None

### 12.1.6 fetch\_run

`health_azure.fetch_run(workspace, run_recovery_id)`

Finds an existing run in an experiment, based on a recovery ID that contains the experiment ID and the actual RunId. The run can be specified either in the `experiment_name:run_id` format, or just the `run_id`.

**Parameters**

- **workspace** (Workspace) – the configured AzureML workspace to search for the experiment.
- **run\_recovery\_id** (str) – The Run to find. Either in the full recovery ID format, `experiment_name:run_id` or just the `run_id`

**Return type** Run

**Returns** The AzureML run.

### 12.1.7 get\_most\_recent\_run

`health_azure.get_most_recent_run(run_recovery_file, workspace)`

Gets the name of the most recently executed AzureML run, instantiates that Run object and returns it.  
 :type run\_recovery\_file: Path :param run\_recovery\_file: The path of the run recovery file :type workspace: Workspace :param workspace: Azure ML Workspace :rtype: Run :return: The Run

### 12.1.8 get\_workspace

`health_azure.get_workspace(aml_workspace, workspace_config_path)`

**Retrieve an Azure ML Workspace from one of several places:**

1. If the function has been called during an AML run (i.e. on an Azure agent), returns the associated workspace
2. If a Workspace object has been provided by the user, return that
3. If a path to a Workspace config file has been provided, load the workspace according to that.

If not running inside AML and neither a workspace nor the config file are provided, the code will try to locate a `config.json` file in any of the parent folders of the current working directory. If that succeeds, that `config.json` file will be used to instantiate the workspace.

**Parameters**

- **aml\_workspace** (Optional[Workspace]) – If provided this is returned as the AzureML Workspace.
- **workspace\_config\_path** (Optional[Path]) – If not provided with an AzureML Workspace, then load one given the information in this config

**Return type** Workspace

**Returns** An AzureML workspace.

### 12.1.9 is\_running\_in\_azure\_ml

`health_azure.is_running_in_azure_ml(aml_run=<azureml.core.run._OfflineRun object>)`

Returns True if the given run is inside of an AzureML machine, or False if it is on a machine outside AzureML. When called without arguments, this functions returns True if the present code is running in AzureML. Note that in runs with “compute\_target=’local’” this function will also return True. Such runs execute outside of AzureML, but are able to log all their metrics, etc to an AzureML run.

**Parameters** `aml_run` (Run) – The run to check. If omitted, use the default run in `RUN_CONTEXT`

**Return type** bool

**Returns** True if the given run is inside of an AzureML machine, or False if it is a machine outside AzureML.

### 12.1.10 set\_environment\_variables\_for\_multi\_node

`health_azure.set_environment_variables_for_multi_node()`

Sets the environment variables that PyTorch Lightning needs for multi-node training.

**Return type** None

### 12.1.11 split\_recovery\_id

`health_azure.split_recovery_id(id)`

Splits a run ID into the experiment name and the actual run. The argument can be in the format ‘experiment\_name:run\_id’, or just a run ID like user\_branch\_abcde12\_123. In the latter case, everything before the last two alphanumeric parts is assumed to be the experiment name.

**Parameters** `id` (str) – The string run ID.

**Return type** Tuple[str, str]

**Returns** experiment name and run name

### 12.1.12 submit\_run

`health_azure.submit_run(workspace, experiment_name, script_run_config, tags=None, wait_for_completion=False, wait_for_completion_show_output=False)`

Starts an AzureML run on a given workspace, via the `script_run_config`.

**Parameters**

- **workspace** (Workspace) – The AzureML workspace to use.
- **experiment\_name** (str) – The name of the experiment that will be used or created. If the experiment name contains characters that are not valid in Azure, those will be removed.
- **script\_run\_config** (Union[ScriptRunConfig, HyperDriveConfig]) – The settings that describe which script should be run.
- **tags** (Optional[Dict[str, str]]) – A dictionary of string key/value pairs, that will be added as metadata to the run. If set to None, a default metadata field will be added that only contains the commandline arguments that started the run.
- **wait\_for\_completion** (bool) – If False (the default) return after the run is submitted to AzureML, otherwise wait for the completion of this run (if True).

- **wait\_for\_completion\_show\_output** (bool) – If wait\_for\_completion is True this parameter indicates whether to show the run output on sys.stdout.

**Return type** Run

**Returns** An AzureML Run object.

### 12.1.13 submit\_to\_azure\_if\_needed

```
health_azure.submit_to_azure_if_needed(compute_cluster_name="", entry_script=None,
                                       aml_workspace=None, workspace_config_file=None,
                                       snapshot_root_directory=None, script_params=None,
                                       conda_environment_file=None, aml_environment_name="",
                                       experiment_name=None, environment_variables=None,
                                       pip_extra_index_url="", private_pip_wheel_path=None,
                                       docker_base_image="", docker_shm_size="",
                                       ignored_folders=None, default_datastore="",
                                       input_datasets=None, output_datasets=None, num_nodes=1,
                                       wait_for_completion=False,
                                       wait_for_completion_show_output=False, max_run_duration="",
                                       submit_to_azureml=None, tags=None, after_submission=None,
                                       hyperdrive_config=None)
```

Submit a folder to Azure, if needed and run it. Use the commandline flag `-azureml` to submit to AzureML, and leave it out to run locally.

#### Parameters

- **after\_submission** (Optional[Callable[[Run], None]]) – A function that will be called directly after submitting the job to AzureML. The only argument to this function is the run that was just submitted. Use this to, for example, add additional tags or print information about the run.
- **tags** (Optional[Dict[str, str]]) – A dictionary of string key/value pairs, that will be added as metadata to the run. If set to None, a default metadata field will be added that only contains the commandline arguments that started the run.
- **aml\_environment\_name** (str) – The name of an AzureML environment that should be used to submit the script. If not provided, an environment will be created from the arguments to this function.
- **max\_run\_duration** (str) – The maximum runtime that is allowed for this job in AzureML. This is given as a floating point number with a string suffix s, m, h, d for seconds, minutes, hours, day. Examples: '3.5h', '2d'
- **experiment\_name** (Optional[str]) – The name of the AzureML experiment in which the run should be submitted. If omitted, this is created based on the name of the current script.
- **entry\_script** (Union[Path, str, None]) – The script that should be run in AzureML
- **compute\_cluster\_name** (str) – The name of the AzureML cluster that should run the job. This can be a cluster with CPU or GPU machines.
- **conda\_environment\_file** (Union[Path, str, None]) – The conda configuration file that describes which packages are necessary for your script to run.
- **aml\_workspace** (Optional[Workspace]) – There are two optional parameters used to glean an existing AzureML Workspace. The simplest is to pass it in as a parameter.

- **workspace\_config\_file** (Union[Path, str, None]) – The 2nd option is to specify the path to the config.json file downloaded from the Azure portal from which we can retrieve the existing Workspace.
- **snapshot\_root\_directory** (Union[Path, str, None]) – The directory that contains all code that should be packaged and sent to AzureML. All Python code that the script uses must be copied over.
- **ignored\_folders** (Optional[List[Union[Path, str]]]) – A list of folders to exclude from the snapshot when copying it to AzureML.
- **script\_params** (Optional[List[str]]) – A list of parameter to pass on to the script as it runs in AzureML. If empty (or None, the default) these will be copied over from sys.argv, omitting the `–azureml` flag.
- **environment\_variables** (Optional[Dict[str, str]]) – The environment variables that should be set when running in AzureML.
- **docker\_base\_image** (str) – The Docker base image that should be used when creating a new Docker image.
- **docker\_shm\_size** (str) – The Docker shared memory size that should be used when creating a new Docker image.
- **pip\_extra\_index\_url** (str) – If provided, use this PIP package index to find additional packages when building the Docker image.
- **private\_pip\_wheel\_path** (Union[Path, str, None]) – If provided, add this wheel as a private package to the AzureML workspace.
- **default\_datastore** (str) – The data store in your AzureML workspace, that points to your training data in blob storage. This is described in more detail in the README.
- **input\_datasets** (Optional[List[Union[str, DatasetConfig]]]) – The script will consume all data in folder in blob storage as the input. The folder must exist in blob storage, in the location that you gave when creating the datastore. Once the script has run, it will also register the data in this folder as an AzureML dataset.
- **output\_datasets** (Optional[List[Union[str, DatasetConfig]]]) – The script will create a temporary folder when running in AzureML, and while the job writes data to that folder, upload it to blob storage, in the data store.
- **num\_nodes** (int) – The number of nodes to use in distributed training on AzureML.
- **wait\_for\_completion** (bool) – If False (the default) return after the run is submitted to AzureML, otherwise wait for the completion of this run (if True).
- **wait\_for\_completion\_show\_output** (bool) – If wait\_for\_completion is True this parameter indicates whether to show the run output on sys.stdout.
- **submit\_to\_azureml** (Optional[bool]) – If True, the codepath to create an AzureML run will be executed. If False, the codepath for local execution (i.e., return immediately) will be executed. If not provided (None), submission to AzureML will be triggered if the command-line flag `–azureml` is present in sys.argv
- **hyperdrive\_config** (Optional[HyperDriveConfig]) – A configuration object for Hyperdrive (hyperparameter search).

**Return type** [AzureRunInfo](#)

**Returns** If the script is submitted to AzureML then we terminate python as the script should be executed in AzureML, otherwise we return a AzureRunInfo object.



### 12.1.14 torch\_barrier

`health_azure.torch_barrier()`

This is a barrier to use in distributed jobs. Use it to make all processes that participate in a distributed pytorch job to wait for each other. When `torch.distributed` is not set up or not found, the function exits immediately.

**Return type** None

### 12.1.15 upload\_to\_datastore

`health_azure.upload_to_datastore(datastore_name, local_data_folder, remote_path, aml_workspace=None, workspace_config_path=None, overwrite=False, show_progress=False)`

Upload a folder to an Azure ML Datastore that is registered within a given Workspace. Note that this will upload all files within the folder, but will not copy the folder itself. E.g. if you specify the `local_data_dir="foo/bar"` and that contains the files ["1.txt", "2.txt"], and you specify the `remote_path="baz"`, you would see the following paths uploaded to your Datastore: ["baz/1.txt", "baz/2.txt"]

If not running inside AML and neither a workspace nor the config file are provided, the code will try to locate a `config.json` file in any of the parent folders of the current working directory. If that succeeds, that `config.json` file will be used to instantiate the workspace.

#### Parameters

- **datastore\_name** (str) – The name of the Datastore to which the blob should be uploaded. This Datastore itself must be an instance of an `AzureBlobDatastore`
- **local\_data\_folder** (Path) – The path to the local directory containing the data to be uploaded
- **remote\_path** (Path) – The path to which the blob should be uploaded
- **aml\_workspace** (Optional[Workspace]) – Optional Azure ML Workspace object
- **workspace\_config\_path** (Optional[Path]) – Optional path to settings for Azure ML Workspace
- **overwrite** (bool) – If True, will overwrite any existing file at the same remote path. If False, will skip any duplicate files and continue to the next.
- **show\_progress** (bool) – If True, will show the progress of the file download

**Return type** None

## 12.2 Classes

<code>AzureRunInfo(input_datasets, ...)</code>	This class stores all information that a script needs to run inside and outside of AzureML.
<code>DatasetConfig(name[, datastore, version, ...])</code>	Contains information to use AzureML datasets as inputs or outputs.

### 12.2.1 AzureRunInfo

```
class health_azure.AzureRunInfo(input_datasets, output_datasets, run, is_running_in_azure_ml,
                                output_folder, logs_folder)
```

Bases: object

This class stores all information that a script needs to run inside and outside of AzureML. It is return from *submit\_to\_azure\_if\_needed*, where the return value depends on whether the script is inside or outside AzureML.

Please check the source code for detailed documentation for all fields.

### 12.2.2 DatasetConfig

```
class health_azure.DatasetConfig(name, datastore="", version=None, use_mounting=None, target_folder="",
                                  local_folder=None)
```

Bases: object

Contains information to use AzureML datasets as inputs or outputs.

#### Parameters

- **name** (str) – The name of the dataset, as it was registered in the AzureML workspace. For output datasets, this will be the name given to the newly created dataset.
- **datastore** (str) – The name of the AzureML datastore that holds the dataset. This can be empty if the AzureML workspace has only a single datastore, or if the default datastore should be used.
- **version** (Optional[int]) – The version of the dataset that should be used. This is only used for input datasets. If the version is not specified, the latest version will be used.
- **use\_mounting** (Optional[bool]) – If True, the dataset will be “mounted”, that is, individual files will be read or written on-demand over the network. If False, the dataset will be fully downloaded before the job starts, respectively fully uploaded at job end for output datasets. Defaults: False (downloading) for datasets that are script inputs, True (mounting) for datasets that are script outputs.
- **target\_folder** (str) – The folder into which the dataset should be downloaded or mounted. If left empty, a random folder on /tmp will be chosen.
- **local\_folder** (Optional[Path]) – The folder on the local machine at which the dataset is available. This is used only for runs outside of AzureML.

#### Methods Summary

<code>to_input_dataset(workspace, dataset_index)</code>	Creates a configuration for using an AzureML dataset inside of an AzureML run.
<code>to_output_dataset(workspace, dataset_index)</code>	Creates a configuration to write a script output to an AzureML dataset.

## Methods Documentation

### **to\_input\_dataset**(*workspace*, *dataset\_index*)

Creates a configuration for using an AzureML dataset inside of an AzureML run. This will make the AzureML dataset with given name available as a named input, using INPUT\_0 as the key for dataset index 0.

#### **Parameters**

- **workspace** (Workspace) – The AzureML workspace to read from.
- **dataset\_index** (int) – Suffix for using datasets as named inputs, the dataset will be marked INPUT\_{index}

**Return type** DatasetConsumptionConfig

### **to\_output\_dataset**(*workspace*, *dataset\_index*)

Creates a configuration to write a script output to an AzureML dataset. The name and datastore of this new dataset will be taken from the present object.

#### **Parameters**

- **workspace** (Workspace) – The AzureML workspace to read from.
- **dataset\_index** (int) – Suffix for using datasets as named inputs, the dataset will be marked OUTPUT\_{index}

**Return type** OutputFileDatasetConfig

#### **Returns**



## HEALTH\_ML.UTILS PACKAGE

### 13.1 Functions

<code>log_learning_rate(module[, name])</code>	Logs the learning rate(s) used by the given module.
<code>log_on_epoch(module[, name, value, metrics, ...])</code>	Write a dictionary with metrics and/or an individual metric as a name/value pair to the loggers of the given module.

#### 13.1.1 log\_learning\_rate

`health_ml.utils.log_learning_rate(module, name='learning_rate')`

Logs the learning rate(s) used by the given module. Multiple learning rate schedulers and/or multiple rates per scheduler are supported. The learning rates are logged under the given name. If multiple scheduler and/or multiple rates are used, a suffix like “/0/1” is added, to indicate the learning rate for scheduler 0, index 1, for example. Learning rates are logged on epoch.

**Parameters**

- **module** (LightningModule) – The module for which the learning rates should be logged.
- **name** (str) – The name to use when logging the learning rates.

**Return type** None

#### 13.1.2 log\_on\_epoch

`health_ml.utils.log_on_epoch(module, name=None, value=None, metrics=None, reduce_fx=<built-in method mean of type object>, sync_dist=None, sync_dist_op='mean')`

Write a dictionary with metrics and/or an individual metric as a name/value pair to the loggers of the given module. Metrics are always logged upon epoch completion. The metrics in question first synchronized across GPUs if DDP with >1 node is used, using the sync\_dist\_op (default: mean). Afterwards, they are aggregated across all steps via the reduce\_fx (default: mean). Metrics that are fed in as plain numbers rather than tensors (for example, plain Python integers) are converted to tensors before logging, to enable synchronization across GPUs if needed.

**Parameters**

- **name** (Optional[str]) – The name of the metric to log.
- **value** (Optional[Any]) – The actual value of the metric to log.
- **metrics** (Optional[Mapping[str, Any]]) – A dictionary with metrics to log.

- **module** (LightningModule) – The PyTorch Lightning module where the metrics should be logged.
- **sync\_dist** (Optional[bool]) – If not None, use this value for the sync\_dist argument to module.log. If None,

set it automatically depending on the use of DDP. Set this to False if you want to log metrics that are only available on Rank 0 of a DDP job. :type reduce\_fx: Callable :param reduce\_fx: The reduce function to apply to the per-step values, after synchronizing the tensors across GPUs. Default: torch.mean :type sync\_dist\_op: Any :param sync\_dist\_op: The reduce operation to use when synchronizing the tensors across GPUs. This must be a value recognized by sync\_ddp: 'sum', 'mean', 'avg'

**Return type** None

## 13.2 Classes

<i>AzureMLLogger()</i>	A Pytorch Lightning logger that stores metrics in the current AzureML run.
------------------------	--

### 13.2.1 AzureMLLogger

**class** health\_ml.utils.AzureMLLogger

Bases: pytorch\_lightning.loggers.base.LightningLoggerBase

A Pytorch Lightning logger that stores metrics in the current AzureML run. If the present run is not inside AzureML, nothing gets logged.

#### Methods Summary

<i>experiment()</i>	Return the experiment object associated with this logger.
<i>log_hyperparams(params)</i>	Record hyperparameters.
<i>log_metrics(metrics[, step])</i>	Records metrics.
<i>name()</i>	Return the experiment name.
<i>version()</i>	Return the experiment version.

#### Methods Documentation

**experiment()**

Return the experiment object associated with this logger.

**Return type** Any

**log\_hyperparams(params)**

Record hyperparameters.

**Parameters**

- **params** (Any) – Namespace containing the hyperparameters
- **args** – Optional positional arguments, depends on the specific logger being used

- **kwargs** – Optional keyword arguments, depends on the specific logger being used

**Return type** None

**log\_metrics**(*metrics*, *step=None*)

Records metrics. This method logs metrics as as soon as it received them. If you want to aggregate metrics for one specific *step*, use the `agg_and_log_metrics()` method.

**Parameters**

- **metrics** (Dict[str, float]) – Dictionary with metric names as keys and measured quantities as values
- **step** (Optional[int]) – Step number at which the metrics should be recorded

**Return type** None

**name()**

Return the experiment name.

**Return type** Any

**version()**

Return the experiment version.

**Return type** int





## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`



## PYTHON MODULE INDEX

### h

`health_azure`, [37](#)

### u

`health_ml.utils`, [49](#)



## INDEX

### A

AzureMLLogger (class in *health\_ml.utils*), 50  
AzureRunInfo (class in *health\_azure*), 46

### C

create\_run\_configuration() (in module *health\_azure*), 38  
create\_script\_run() (in module *health\_azure*), 39

### D

DatasetConfig (class in *health\_azure*), 46  
download\_checkpoints\_from\_run\_id() (in module *health\_azure*), 40  
download\_files\_from\_run\_id() (in module *health\_azure*), 39  
download\_from\_datastore() (in module *health\_azure*), 40

### E

experiment() (*health\_ml.utils.AzureMLLogger* method), 50

### F

fetch\_run() (in module *health\_azure*), 41

### G

get\_most\_recent\_run() (in module *health\_azure*), 41  
get\_workspace() (in module *health\_azure*), 41

### H

health\_azure  
    module, 37  
health\_ml.utils  
    module, 49

### I

is\_running\_in\_azure\_ml() (in module *health\_azure*), 42

### L

log\_hyperparams() (*health\_ml.utils.AzureMLLogger* method), 50

log\_learning\_rate() (in module *health\_ml.utils*), 49  
log\_metrics() (*health\_ml.utils.AzureMLLogger* method), 51  
log\_on\_epoch() (in module *health\_ml.utils*), 49

### M

module  
    *health\_azure*, 37  
    *health\_ml.utils*, 49

### N

name() (*health\_ml.utils.AzureMLLogger* method), 51

### S

set\_environment\_variables\_for\_multi\_node() (in module *health\_azure*), 42  
split\_recovery\_id() (in module *health\_azure*), 42  
submit\_run() (in module *health\_azure*), 42  
submit\_to\_azure\_if\_needed() (in module *health\_azure*), 43

### T

to\_input\_dataset() (*health\_azure.DatasetConfig* method), 47  
to\_output\_dataset() (*health\_azure.DatasetConfig* method), 47  
torch\_barrier() (in module *health\_azure*), 45

### U

upload\_to\_datastore() (in module *health\_azure*), 45

### V

version() (*health\_ml.utils.AzureMLLogger* method), 51