

Distributed Deep Learning Using Hopsworks

CGI Trainee Program Workshop

Kim Hammar

kim@logicalclocks.com



LOGICAL CLOCKS

Before we start..

1. Register for an account at:


www.hops.site

2. Follow the instructions at:


<http://bit.ly/2EnZQgW>

DISTRIBUTED COMPUTING + DEEP LEARNING = ?

Distributed Computing



Deep Learning




Why Combine the two?

2em1¹ Chen Sun et al. “Revisiting Unreasonable Effectiveness of Data in Deep Learning Era”. In: *CoRR* abs/1707.02968 (2017). arXiv: [1707.02968](https://arxiv.org/abs/1707.02968). URL: <http://arxiv.org/abs/1707.02968>.


2em1² Jeffrey Dean et al. “Large Scale Distributed Deep Networks”. In: *Advances in Neural Information Processing Systems* 25. Ed. by F. Pereira et al. Curran Associates, Inc., 2012, pp. 1223–1231. ◀ ◻ ▶ ◀ ◻ ▶ ◀ ≡ ▶ ◀ ≡ ▶ ≡ 🔍 ↻

DISTRIBUTED COMPUTING + DEEP LEARNING = ?

Distributed Computing



Deep Learning



Why Combine the two?


- ▶ We like challenging problems 😊

2em1¹ Chen Sun et al. "Revisiting Unreasonable Effectiveness of Data in Deep Learning Era". In: *CoRR* abs/1707.02968 (2017). arXiv: 1707.02968. URL: <http://arxiv.org/abs/1707.02968>.

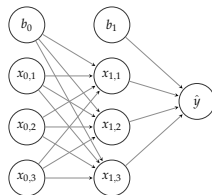
2em1² Jeffrey Dean et al. "Large Scale Distributed Deep Networks". In: *Advances in Neural Information Processing Systems* 25. Ed. by F. Pereira et al. Curran Associates, Inc., 2012, pp. 1223–1231. ◀ ◻ ▶ ◀ ☰ ▶ ◀ ☷ ▶ ◀ ☹ ▶ ◀ ☰ ▶ ◀ ☷ ▶ ◀ ☹ ▶

DISTRIBUTED COMPUTING + DEEP LEARNING = ?

Distributed Computing



Deep Learning




Why Combine the two?

- ▶ We like challenging problems 😊
- ▶ More productive data science
- ▶ Unreasonable effectiveness of data¹
- ▶ To achieve state-of-the-art results²

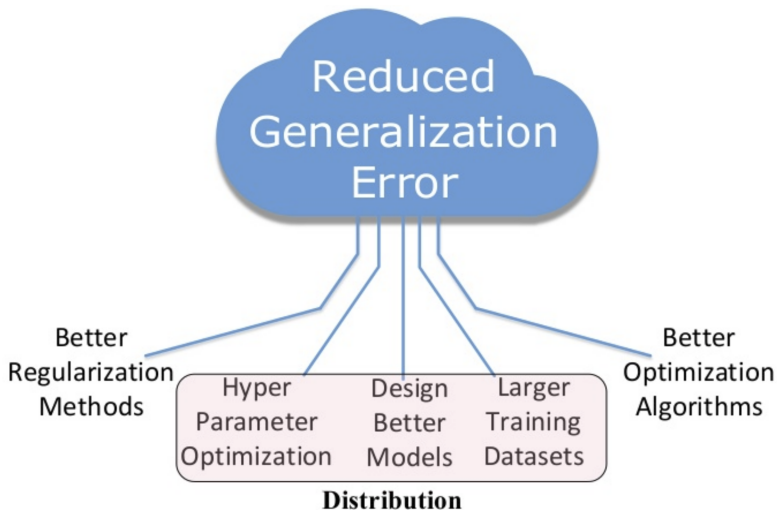
2em1¹ Chen Sun et al. “Revisiting Unreasonable Effectiveness of Data in Deep Learning Era”. In: *CoRR* abs/1707.02968 (2017). arXiv: [1707.02968](https://arxiv.org/abs/1707.02968). URL: <http://arxiv.org/abs/1707.02968>.

2em1² Jeffrey Dean et al. “Large Scale Distributed Deep Networks”. In: *Advances in Neural Information Processing Systems* 25. Ed. by F. Pereira et al. Curran Associates, Inc., 2012, pp. 1223–1231. ◀ ◻ ▶ ◀ ◻ ▶ ◀ ≡ ▶ ◀ ≡ ▶ ≡ 🔍 ↻


DISTRIBUTED DEEP LEARNING (DDL): PREDICTABLE SCALING



DISTRIBUTED DEEP LEARNING (DDL): PREDICTABLE SCALING



DDL IS NOT A SECRET ANYMORE



(b) Training with Single vs. Multiple Nodes

4

DDL IS NOT A SECRET ANYMORE

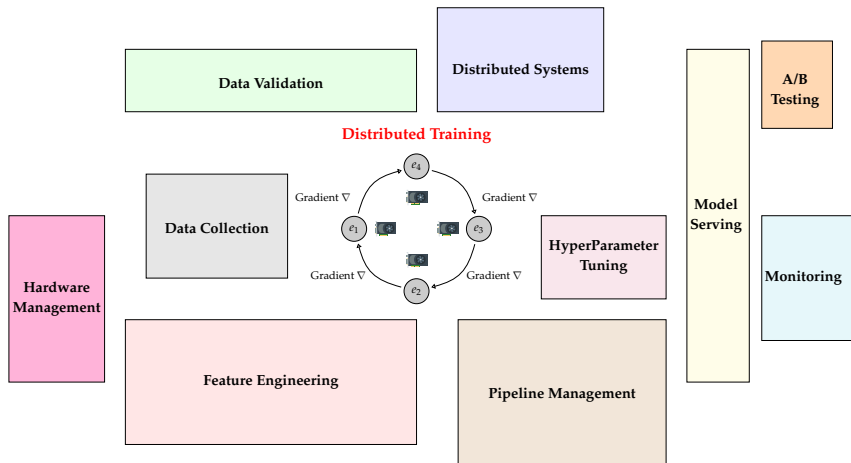
Frameworks for DDL



Companies using DDL



DDL REQUIRES AN ENTIRE SOFTWARE/INFRASTRUCTURE STACK




OUTLINE

1. **Hopsworks**: Background of the platform
2. **Managed Distributed Deep Learning** using HopsYARN, HopsML, PySpark, and Tensorflow
3. **Black-Box Optimization (Hyperparameter Tuning)** using Hopsworks, Metadata Store, PySpark, and **Maggy**⁵
4. **Feature Store** data management for machine learning
5. **Coffee Break**
6. **Demo**, end-to-end ML pipeline
7. **Hands-on Workshop**, try out Hopsworks on our cluster in Luleå

HOPSWORKS

HOPSWORKS




HOPSWORKS

HopsYARN

(GPU/CPU as a resource)



HopsFS



HOPSWORKS

Frameworks

(ML/Data)



PYTORCH






HopsYARN

(GPU/CPU as a resource)




HopsFS



HOPSWORKS

ML/AI Assets



Frameworks

(ML/Data)




HopsYARN

(GPU/CPU as a resource)



HopsFS



HOPSWORKS


APIs

```
from hops import featurestore
from hops import experiment
featurestore.get_features([
    "average_attendance",
    "average_player_age"])
experiment.collective_all_reduce(features, model)
```

ML/AI Assets



Feature Store



Pipelines



Experiments



Models

Frameworks

(ML/Data)



PYTORCH



HopsYARN

(GPU/CPU as a resource)



HopsFS



HOPSWORKS


APIs

```
from hops import featurestore
from hops import experiment
featurestore.get_features([
    "average_attendance",
    "average_player_age"])
experiment.collective_all_reduce(features, model)
```

ML/AI Assets



Feature Store



Pipelines



Experiments



Models

Frameworks

(ML/Data)



PYTORCH




HopsYARN

(GPU/CPU as a resource)




Distributed Metadata

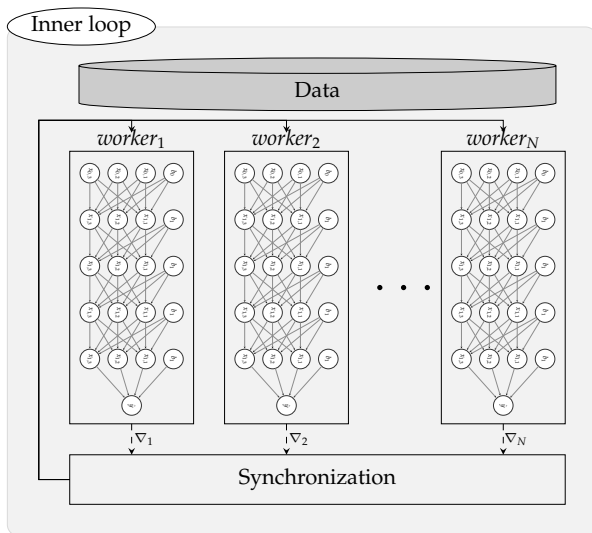
(Available from REST API)




HopsFS



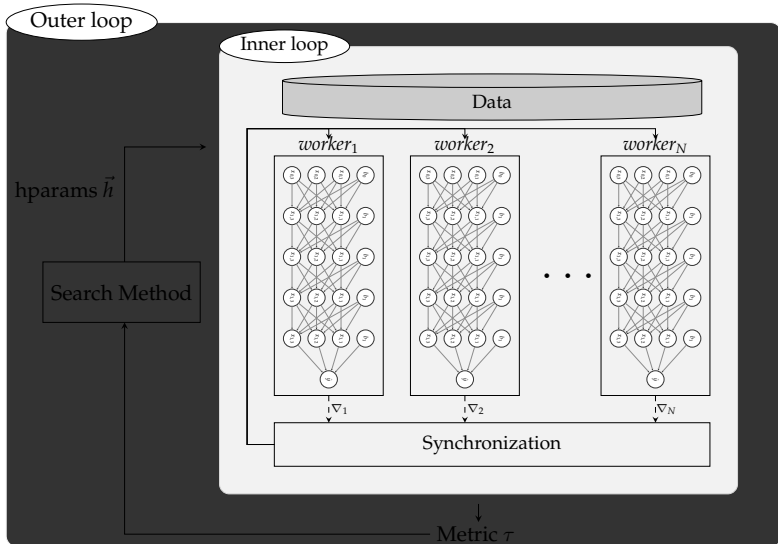
INNER AND OUTER LOOP OF LARGE SCALE DEEP LEARNING



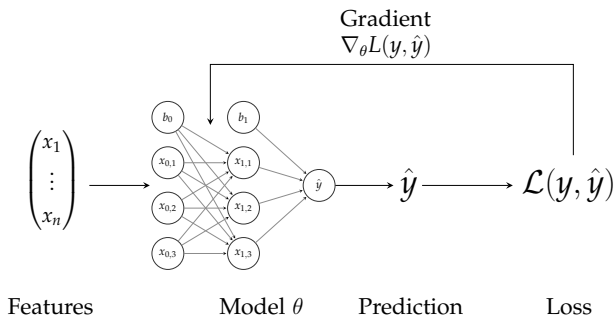
INNER AND OUTER LOOP OF LARGE SCALE DEEP LEARNING



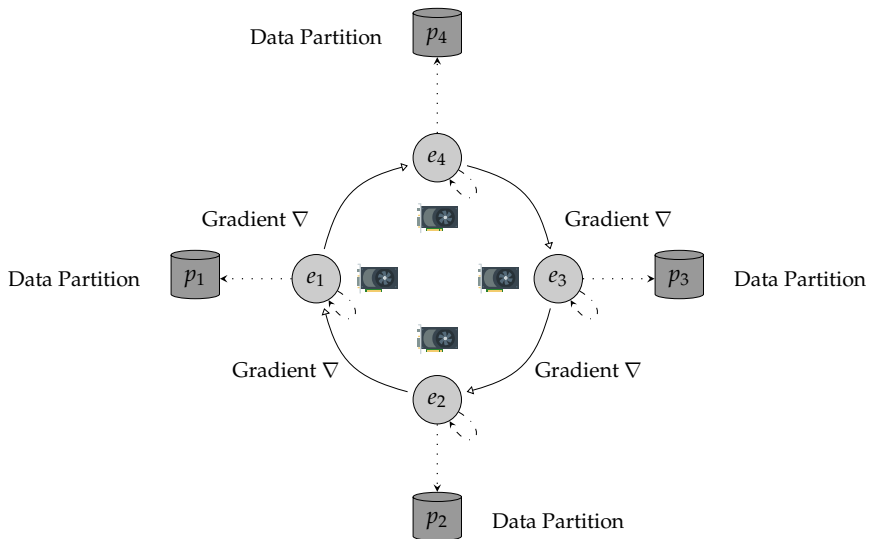
INNER AND OUTER LOOP OF LARGE SCALE DEEP LEARNING



INNER LOOP: DISTRIBUTED DEEP LEARNING




INNER LOOP: DISTRIBUTED DEEP LEARNING



DISTRIBUTED DEEP LEARNING IN PRACTICE

- ▶ Implementation of distributed algorithms is becoming a **commodity** (TF, PyTorch etc)
- ▶ **The hardest part of DDL is now:**
 - ▶ Cluster management
 - ▶ Allocating GPUs
 - ▶ Data management
 - ▶ Operations & performance



HOPSWORKS DDL SOLUTION

HOPSWORKS DDL SOLUTION

```
from hops import experiment
experiment.collective_all_reduce(train_fn)
```

HOPSWORKS DDL SOLUTION

```

from hops
import experiment
experiment.
collective_all_reduce(
train_fn
)

```

Client API



HOPSWORKS DDL SOLUTION

```
from hops
import experiment
experiment.
collective_all_reduce(
train_fn
)
```

Client API



HOPSWORKS DDL SOLUTION





```

from hops
import experiment
experiment.
collective_all_reduce(
train_fn
)

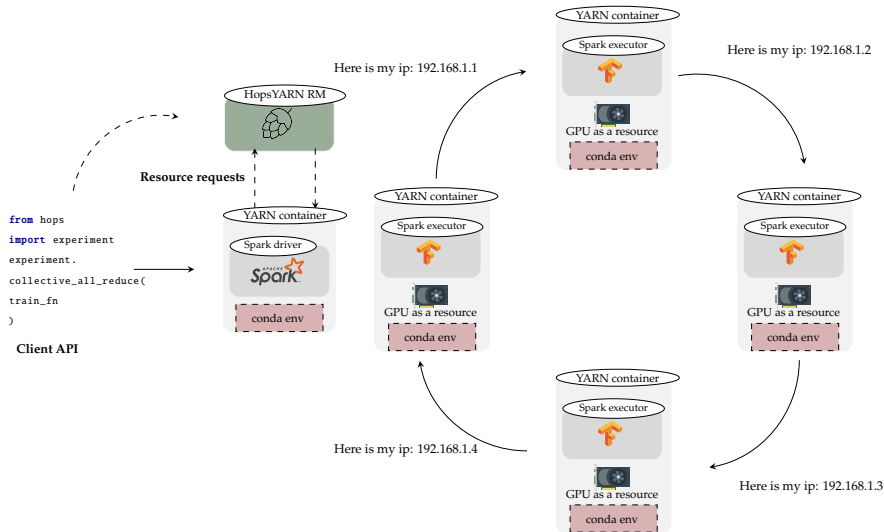
```

Client API

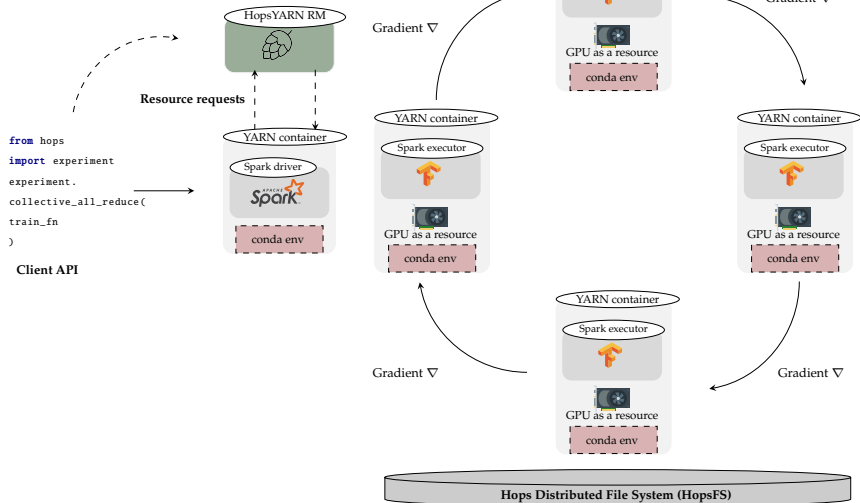
Resource requests



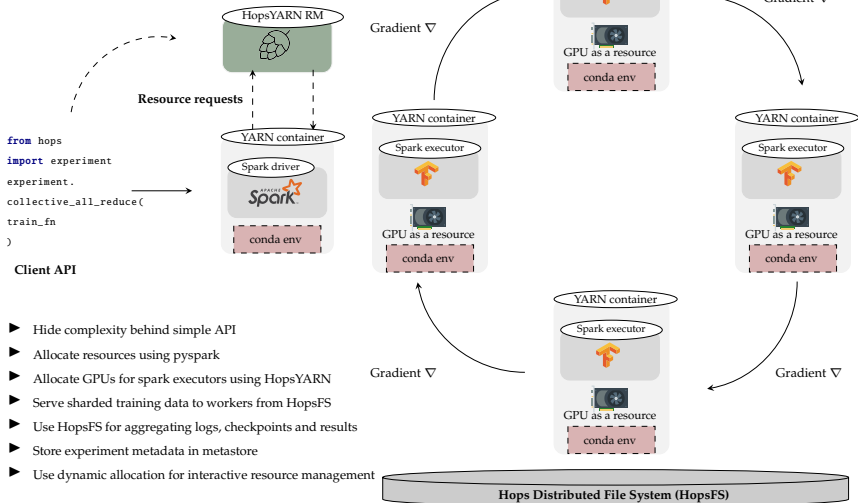
HOPSWORKS DDL SOLUTION




HOPSWORKS DDL SOLUTION




HOPSWORKS DDL SOLUTION



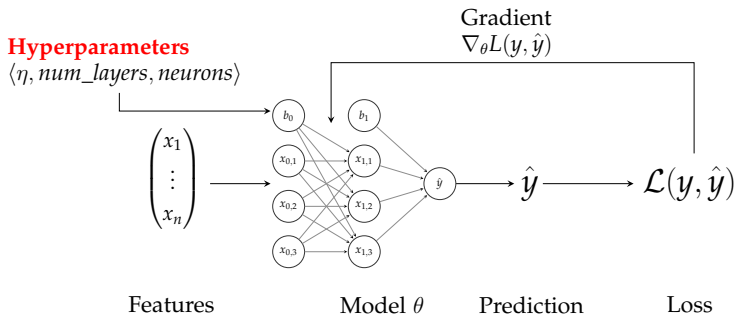
OUTER LOOP: BLACK BOX OPTIMIZATION



OUTER LOOP: BLACK BOX OPTIMIZATION




OUTER LOOP: BLACK BOX OPTIMIZATION



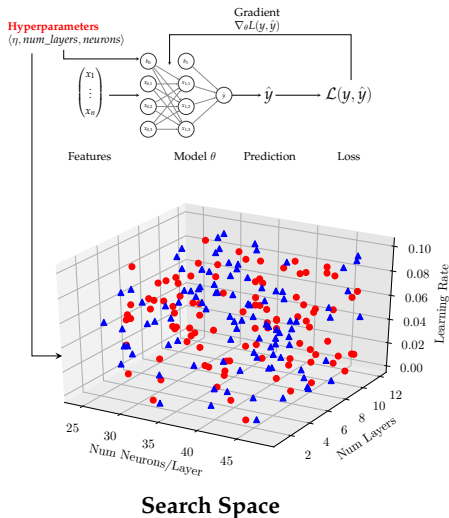
OUTER LOOP: BLACK BOX OPTIMIZATION

Example Use-Case from one of our clients:



- ▶ **Goal: Train a One-Class GAN model for fraud detection**
- ▶ Problem: GANs are extremely sensitive to hyperparameters and there exists a very large space of possible hyperparameters.
- ▶ Example hyperparameters to tune: learning rates η , optimizers, layers.. etc.




OUTER LOOP: BLACK BOX OPTIMIZATION



OUTER LOOP: BLACK BOX OPTIMIZATION




Search Space





Shared Task Queue


Parallel Workers



OUTER LOOP: BLACK BOX OPTIMIZATION




Search Space




Shared Task Queue

Parallel Workers

OUTER LOOP: BLACK BOX OPTIMIZATION




Which algorithm to use for search?



Search Space




Shared Task Queue




Parallel Workers

OUTER LOOP: BLACK BOX OPTIMIZATION




How to monitor progress?

Which algorithm to use for search?



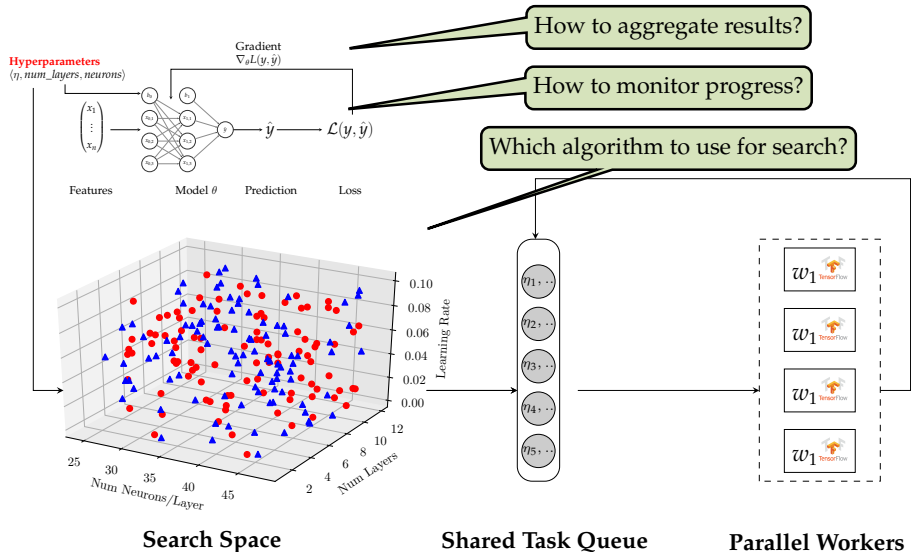
Search Space



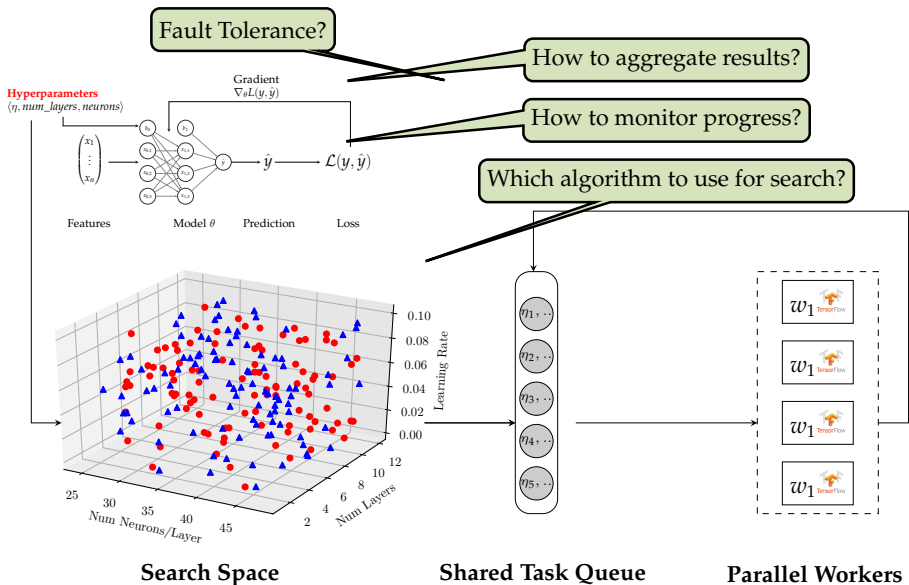
Shared Task Queue

Parallel Workers

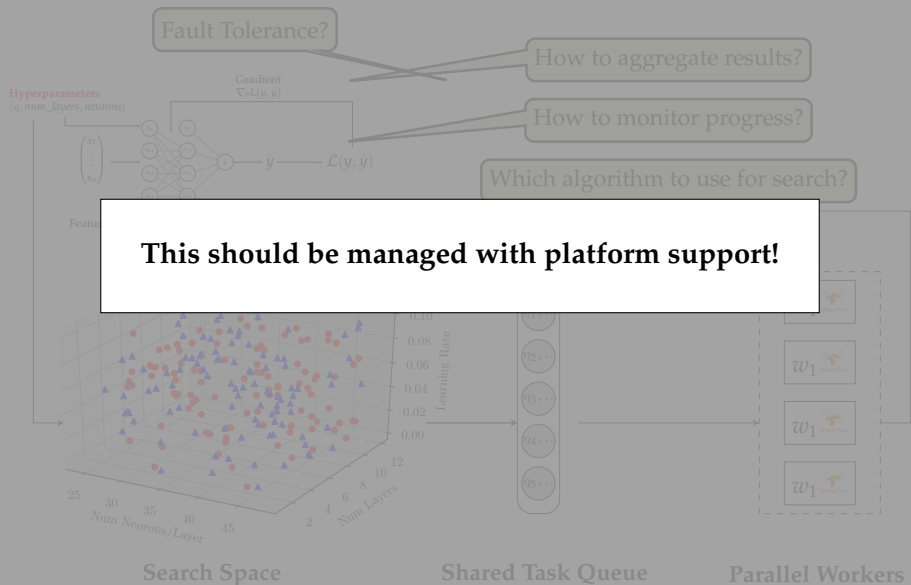
OUTER LOOP: BLACK BOX OPTIMIZATION



OUTER LOOP: BLACK BOX OPTIMIZATION



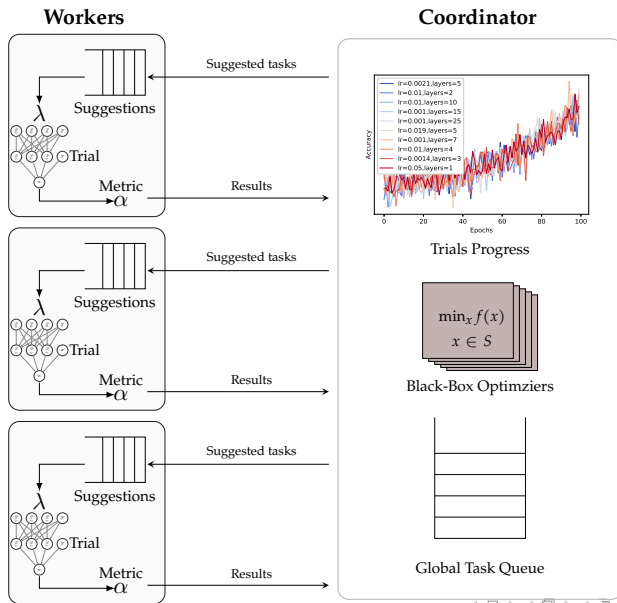
OUTER LOOP: BLACK BOX OPTIMIZATION



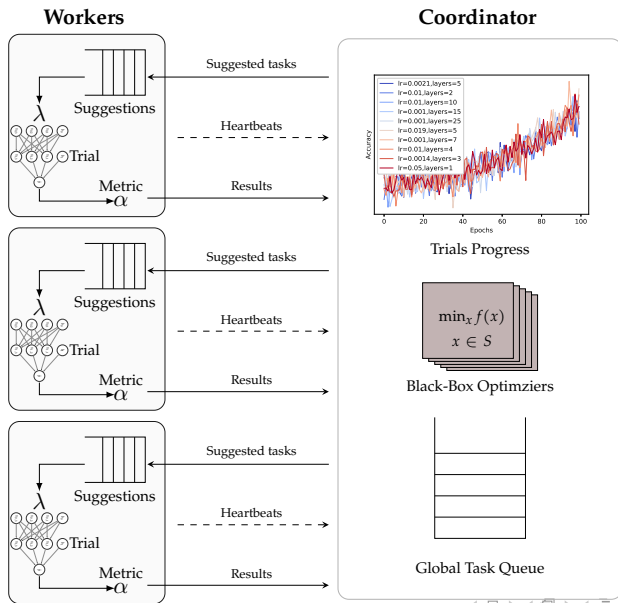
PARALLEL EXPERIMENTS

```
from hops import experiment
experiment.random_search(train_fn)
```


ASYNCHRONOUS SEARCH WORKFLOW



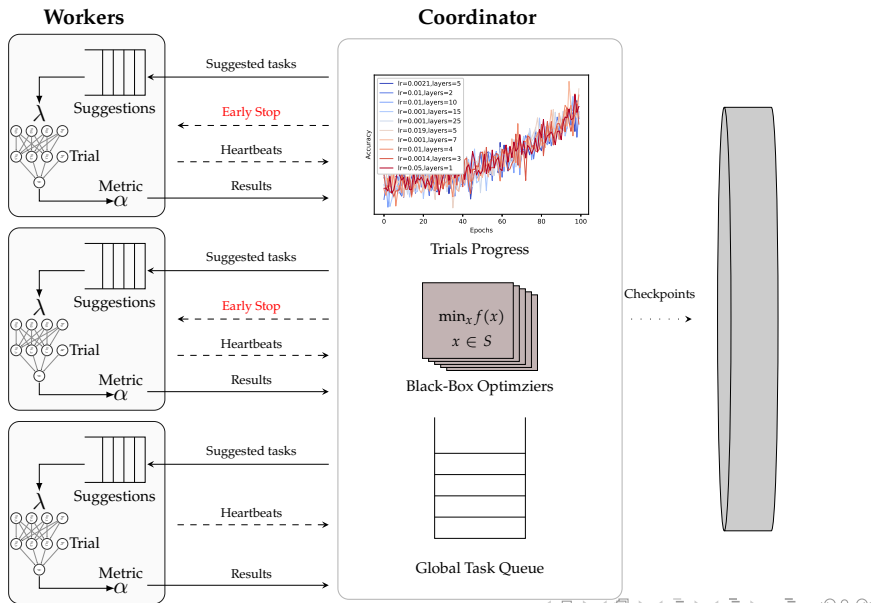
ASYNCHRONOUS SEARCH WORKFLOW




ASYNCHRONOUS SEARCH WORKFLOW



ASYNCHRONOUS SEARCH WORKFLOW




INNER AND OUTER LOOP OF LARGE SCALE DEEP LEARNING




FEATURE STORE


FEATURE STORE



FEATURE STORE



FEATURE STORE




“Data is the hardest part of ML and the most important piece to get right.”

Modelers spend most of their time selecting and transforming features at training time and then building the pipelines to deliver those features to production models.”

- Uber⁶

FEATURE STORE



"Data is the hardest part of ML and the most important piece to get right."

Modelers spend most of their time selecting and transforming features at training time and then building the pipelines to deliver those features to production models."

- Uber⁷

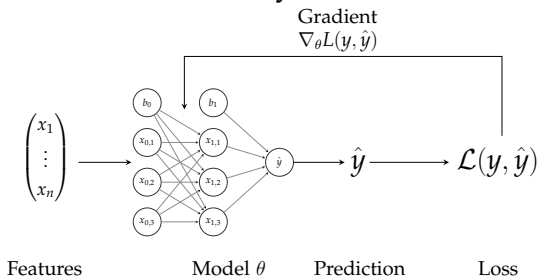
WHAT IS A FEATURE?

A feature is a measurable property of some data-sample


A feature could be..

- ▶ An aggregate value (min, max, mean, sum)
- ▶ A raw value (a pixel, a word from a piece of text)
- ▶ A value from a database table (the age of a customer)
- ▶ A derived representation: e.g an embedding or a cluster


Features are the fuel for AI systems:




FEATURE ENGINEERING IS CRUCIAL FOR MODEL PERFORMANCE



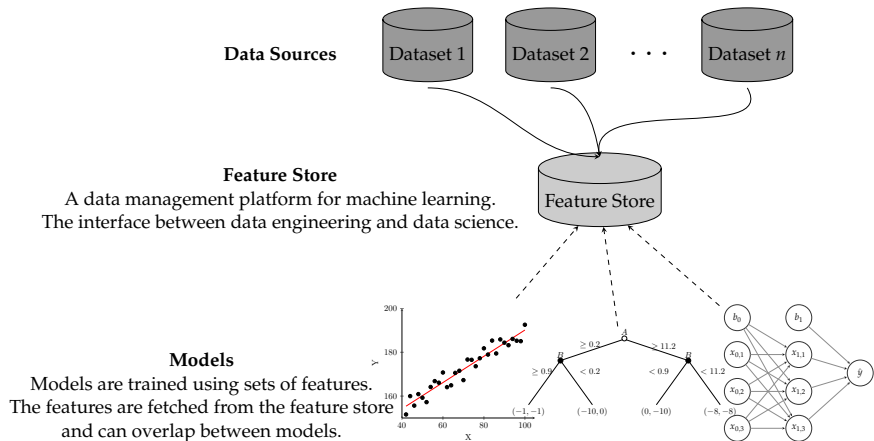
FEATURE ENGINEERING IS CRUCIAL FOR MODEL PERFORMANCE



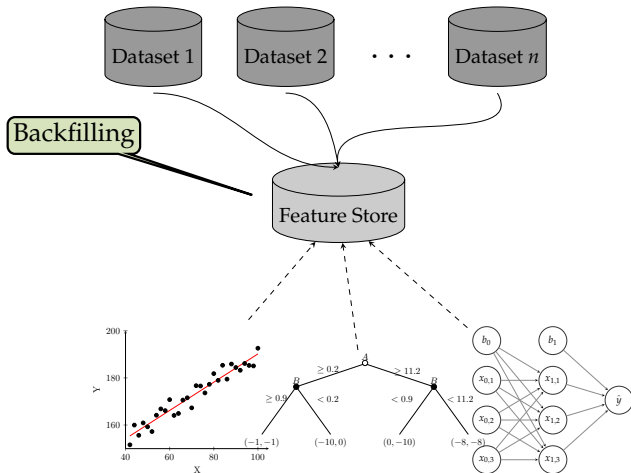
FEATURE ENGINEERING IS CRUCIAL FOR MODEL PERFORMANCE



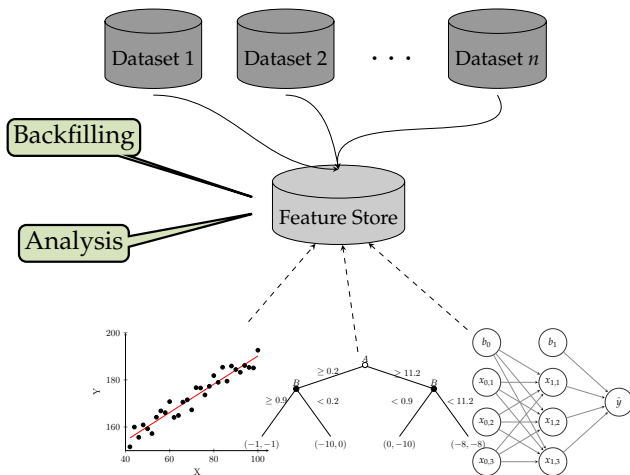
DISENTANGLE YOUR ML PIPELINES WITH A FEATURE STORE



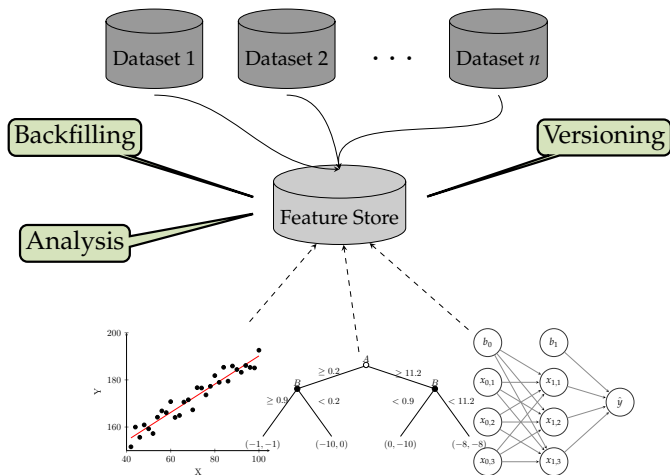
DISENTANGLE YOUR ML PIPELINES WITH A FEATURE STORE



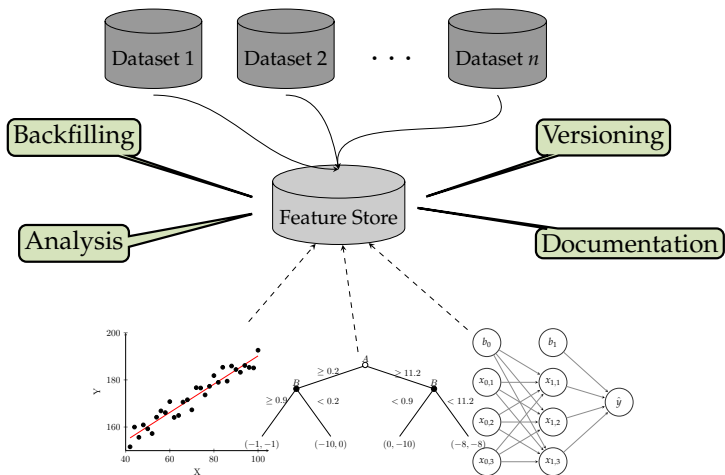
DISENTANGLE YOUR ML PIPELINES WITH A FEATURE STORE



DISENTANGLE YOUR ML PIPELINES WITH A FEATURE STORE



DISENTANGLE YOUR ML PIPELINES WITH A FEATURE STORE



SUMMARY

- ▶ Deep Learning is going distributed
- ▶ Algorithms for DDL are available in several frameworks
- ▶ Applying DDL in practice brings a lot of operational complexity
- ▶ Hopsworks is a platform for scale out deep learning and big data processing
- ▶ Hopsworks makes DDL simpler by providing simple abstractions for distributed training, parallel experiments and much more..



@hopshadoop

www.hops.io

@logicalclocks



www.logicalclocks.com

LOGICAL CLOCKS

We are open source:

<https://github.com/logicalclocks/hopsworks>
<https://github.com/hopshadoop/hops>

Thanks to Logical Clocks Team: Jim Dowling, Seif Haridi, Theo Kakantousis, Fabio Buso, Gautier Berthou, Ermias Gebremeskel, Mahmoud Ismail, Salman Niazi, Antonios Kouzoupis, Robin Andersson, Alex Ormenisan, Rasmus Toivonen and Steffen Grohsschmiedt.

Demo-Setting

Raw/Structured Data




Feature
Computation



Curated Features



Model



Hands-on Workshop

1. If you haven't registered, do it now on hops.site
2. Cheatsheet: http://snurran.sics.se/hops/kim/workshop_cheat.txt

EXERCISE 1 (HELLO HOPSWORKS)

1. Create a Deep Learning Tour Project on Hopsworks
2. Start a Jupyter Notebook with the config:
 - ▶ “Experiment” Mode
 - ▶ 1 GPU
 - ▶ 4000 (MB) memory for the driver (appmaster)
 - ▶ 8000 (MB) memory for the executor
 - ▶ Rest can be default
3. Create a new “PySpark” notebook
4. In the first cell, write:

```
print("Hello Hopsworks")
```
5. Execute the cell (Ctrl + <Enter>)

EXERCISE 2 (DISTRIBUTED HELLO HOPSWORKS WITH GPU)

1. Add a new cell with the contents:


```
def executor():  
    print("Hello from GPU")
```

2. Add a new cell with the contents:

```
from hops import experiment  
experiment.launch(executor)
```

3. Execute the two cells in order (Ctrl + <Enter>)
4. Go to the Application UI

EXERCISE 2 (DISTRIBUTED HELLO HOPSWORKS WITH GPU)



EXERCISE 2 (DISTRIBUTED HELLO HOPSWORKS WITH GPU)

```
In [ ]:
```

```
def executor():  
    print("Hello from GPU")
```

```
In [ ]:
```

```
def executor():  
    print("Hello from GPU")
```

```
In [ ]:
```

```
def executor():  
    print("Hello from GPU")
```

```
In [ ]:
```

```
def executor():  
    print("Hello from GPU")
```

```
In [ ]:
```

```
from hops import experiment  
experiment.launch(executor)
```


EXERCISE 3 (LOAD MNIST FROM HOPSF)

1. Add a new cell with the contents:

```
from hops import hdfs
import tensorflow as tf
def create_tf_dataset():
    train_files = [hdfs.project_path() +
                   "TestJob/data/mnist/train/train.tfrecords"]
    dataset = tf.data.TFRecordDataset(train_files)
def decode(example):
    example = tf.parse_single_example(example, {
        'image_raw': tf.FixedLenFeature([], tf.string),
        'label': tf.FixedLenFeature([], tf.int64)})
    image = tf.reshape(tf.decode_raw(example['image_raw'],
                                     tf.uint8), (28,28,1))
    label = tf.one_hot(tf.cast(example['label'], tf.int32), 10)
    return image, label
return dataset.map(decode).batch(128).repeat()
```

EXERCISE 3 (LOAD MNIST FROM HOPSF5)

1. Add a new cell with the contents:

```
create_tf_dataset()
```

2. Execute the two cells in order (Ctrl + <Enter>)

EXERCISE 4 (DEFINE CNN MODEL)

```
from tensorflow import keras
def create_model():
    model = keras.Sequential()
    model.add(keras.layers.Conv2D(filters=32, kernel_size=3, padding='same',
                                   activation='relu', input_shape=(28,28,1)))
    model.add(keras.layers.BatchNormalization())
    model.add(keras.layers.MaxPooling2D(pool_size=2))
    model.add(keras.layers.Dropout(0.3))
    model.add(keras.layers.Conv2D(filters=64, kernel_size=3,
                                   padding='same', activation='relu'))
    model.add(keras.layers.BatchNormalization())
    model.add(keras.layers.MaxPooling2D(pool_size=2))
    model.add(keras.layers.Dropout(0.3))
    model.add(keras.layers.Flatten())
    model.add(keras.layers.Dense(128, activation='relu'))
    model.add(keras.layers.Dropout(0.5))
    model.add(keras.layers.Dense(10, activation='softmax'))
    return model
```

EXERCISE 4 (DEFINE CNN MODEL)

1. Add a new cell with the contents:

```
create_model().summary()
```

2. Execute the two cells in order (Ctrl + <Enter>)

EXERCISE 5 (DEFINE & RUN THE EXPERIMENT)

1. Add a new cell with the contents:

```
from hops import tensorboard
from tensorflow.python.keras.callbacks import TensorBoard
def train_fn():
    dataset = create_tf_dataset()
    model = create_model()
    model.compile(loss=keras.losses.categorical_crossentropy,
                  optimizer=keras.optimizers.Adam(), metrics=['accuracy'])
    tb_callback = TensorBoard(log_dir=tensorboard.logdir())
    model_checkpoint_callback = keras.callbacks.ModelCheckpoint(
        tensorboard.logdir(), monitor='acc')
    history = model.fit(dataset, epochs=50,
                        steps_per_epoch=80, callbacks=[tb_callback])
    return history.history["acc"][-1]
```

EXERCISE 5 (DEFINE & RUN THE EXPERIMENT)

1. Add a new cell with the contents:

```
experiment.launch(train_fn)
```

2. Execute the two cells in order (Ctrl + <Enter>)
3. Go to the Application UI and monitor the training progress

REFERENCES

- ▶ Example notebooks <https://github.com/logicalclocks/hops-examples>
- ▶ HopsML⁸
- ▶ Hopsworks⁹
- ▶ Hopsworks' feature store¹⁰
- ▶ Maggy
<https://github.com/logicalclocks/maggy>

2em1⁸ Logical Clocks AB. *HopsML: Python-First ML Pipelines*. <https://hops.readthedocs.io/en/latest/hopsml/hopsML.html>. 2018.

2em1⁹ Jim Dowling. *Introducing Hopsworks*. <https://www.logicalclocks.com/introducing-hopsworks/>. 2018.

2em1¹⁰ Kim Hammar and Jim Dowling. *Feature Store: the missing data layer in ML pipelines?* <https://www.logicalclocks.com/feature-store/>. 2018.