

Deep Learning with Databricks

Day 1:

- Neural network and tf.keras fundamentals
- Improve models by adding data standardization, callbacks, checkpointing, etc.
- Track and version models with MLflow
- Distributed inference with pandas UDFs
- Distributed hyper parameter tuning with Hyper opt
- Large scale data preparation with Petastorm

Day 2:

- Distributed model training with Horovod and Petastorm
- Model interpretability with SHAP
- CNNs for image classification and transfer learning
- Distributed training with TFRecord using spark-tensor flow-distributor
- Deploy REST endpoint using MLflow Model Serving on Databricks
- Textual embedding, RNNs, attention-based models, and transfer learning for named entity recognition (NER)