



Architect's Guide for Continuous Machine Learning Platforms With Apache Ignite 2.8

Ken Cottrell Solution Architect ken.cottrell@gridgain.com





Architect's Guide for Continuous Machine Learning With Apache Ignite 2.8

- Introduction to ML/DL Architecture Patterns: Continuous Machine Learning at Scale
- Apache Ignite 2.8 ML/DL Capabilities and New Features
- Apache Ignite 2.8 ML/DL Code Examples / Demos
- Q & A



Apache Ignite 2.8



Over 1,900 upgrades and fixes that enhance almost all platform components

- Production Support Improvements to better support high 24x7 workloads
 - Persistence compaction, Baseline topology autoadjustment
- Support for Next-Generation Monitoring, User-defined Alerts, and Event Tracing
 - APIs for JMX, SQL, OpenCensus
 - Take a look at the new Gridgain Control Center product: https://control.gridgain.com/
- Machine Learning Enhancements
 - A new pipelining API gives you added ML workflow flexibility
 - Additional <u>ensemble methods</u>, which allows you to combine ML training algorithms into more agile training workflows
 - Model parsers to leverage outputs from such leading ML technologies as <u>Apache Spark and</u> <u>XGBoost models</u>



Apache Ignite Is a Top 5 Apache Project





From January 1, 2019 Apache Software Foundation Blog Post: "Apache in 2018 – By The Digits"

GridGai

2019 © GridGain Systems

Apache Ignite Users









2019 © GridGain Systems

Apache Ignite In-Memory Computing Platform



Apache Ignite Features

GridGain Enterprise Features



ML/DL Architectures that deliver Continuous Machine Learning at Scale





Continuous Machine Learning At Scale: Both Real-time and Batch Data, Together

- New "model update" APIs let refine an existing model on the fly with new data samples
- Ignite ML is deeply integrated into Ignite multi-tiered storage:
 - Optional ETL and no data shuffling during training
 - (Re-)Train across petabytes of inmemory and on-disk data





Continuous Machine Learning At Scale: Enabled with Partitioned Datasets





Continuous ML at Scale: Example Deployment Patterns







Deployment Pattern: Building Predictive Models





Batch training on large datasets

Latency constraint:

- Batch run
- Approx. 1 / Node count time reduction

Pattern: ML tasks performed where OLTP data already resides

- Extract OLTP into ML dataset
- Preprocess into ML format
- Train w/ selected algorithms
- Evaluate for accuracy

Output:

- Trained ML predictive Model(s) collected in client node(s)
- Select models using evaluation data and other criteria
- Models can be saved and loaded back to other clients



Deployment Pattern: Using the Predictive Models Real-Time predictions





Latency constraint:

 Per OLTP commit: Application needs to get prediction first before performing commit

Pattern:

- Predictive Model forward deployed close to OLTP application for fast response
- Options include: Same JVM, Same Node, Same cluster

- OLTP writes performed in 1/N scaleout
- OLTP writes include both Predictions + Actuals for accuracy tracking



Deployment Pattern: Using the Predictive Models "Near Real-Time" predictions



Latency constraint:

- Prediction doesn't need to be sent back to application client before write to cache; predictions can be done in reasonable size batches
- Workflow prioritization post commit: Claims processing for example

Pattern:

- Predictive Model forward deployed close to OLTP application
- Options include: Same JVM, Same Node, Same cluster

- OLTP writes performed in 1/N scale-out
- OLTP writes include both Predictions +
 Actuals for accuracy tracking at each node



Deployment Pattern: Using the Predictive Models Shared Model as a service



Pattern:

- Replicated OLTP client nodes that share Model service
- Predictive Model forward deployed close to OLTP application, if you need you need predicts prior to commits
- Wrap service for governance (security, metering, routing / transformation)

- OLTP writes performed in 1/N scale-out
- OLTP writes include both Predictions + Actuals for accuracy tracking



Deployment Pattern: Using the Predictive Models Replicating the Model service



Latency constraint:

 Can supply prediction to application prior to OLTP commit

Pattern:

- Predictive Model forward deployed but now load-balanced
- Identical Model is replicated (Deserialized or Parsed) into each Model node.
- Model needs to be loaded from one source so identical in each instance

- OLTP writes performed in 1/N scale-out
- OLTP writes include both Predictions + Actuals for accuracy tracking



Deployment Pattern: Using the Predictive Models Batch process





Latency constraints

- orchestrating repeat batch runs: no per-record round trips to Application on the way to cluster
- Not always a "predict" problem perse since you might be looking for unknowns
- Examples: clustering for unlabeled population health, customers, network behavior

Deployment Pattern

 Forward deployed Model or colocated Model with data, depends: May require interactive 3rd party AutoML orchestrations via API

Outputs / Results

- Model results collected on client (i.e. Entropy evaluation if Clustering) Parallel processing with co-located data
- Extract, preprocess, train, evaluate, rerun with new data, training, evaluation, reporting.....



-

15

Apache Ignite 2.8 ML/DL

Pipeline APIs, Algorithms





Apache Ignite Continuous Learning framework





Apache Ignite 2.8 ML Out-Of-The-Box



Algorithms for every ML Pipeline use case

Pipeline Stage	Output
Data Extraction Preprocessing	 Feature Vectorization (i.e. columns of interest in ML training) Handling the major column types: string, categoricals, numbers Normalization of variables Preventing overfitting (Cross Validate, Test/Train Split)
Training Evaluation (before deployment)	 Regression labeled data for numeric estimation Classifying labeled data for predicting output classes Importing from 3rd party work benches XGBoost, Spark, H20 Loading previously trained & saved Ignite models
Prediction Performance Retraining	 Extensive set of predictive Algorithms including newly released Ensembles Transaction & Batch workloads Deployment flexibility: Either front-deployed near your Application client or co-located with OLTP cache data



Apache Ignite Distributed PreProcessing: Normalization & Scaling

Improved Model training velocity and accuracy

- Prevent large magnitude variables ٠ from overwhelming smaller ones
- Speed up processing by re-centering



https://medium.com/@nsethi610/data-cleaning-scale-and-normalize-data-4a7c781dd628



Apache Ignite Distributed preprocessing: Categorical values: One-Hot encoders



* Also included: String Encoding



Apache Ignite Distributed Training: Clustering



For "unlabeled" data – the data appears continuous but are there actually classes like "high risk", "medium risk", "low risk"?

Examples:

- Population Segmentation (health, retail, genomics)
- Anomaly Detection
- Network throughput characterization



K-means, GMM



Apache Ignite Distributed Training: Classification & Regression

Predicting output classes

- Fraud detection
- Credit Card Scoring
- Clinical Trials
- Customer Segmentation

Numerical Estimation

- Trend Analysis
- Financial Forecasting
- Time Series Prediction
- Response Modeling (Pharma, etc)





- Logistic Regression & Naive Bayes
- SVM, KNN, ANN
- Decision trees & Random Forest
- Import XGBoost, Spark
- KNN & Linear Regressions
- Decision tree regression
- Random forest regression
- Gradient-boosted tree regression
- Import XGBoost, Spark





Apache Ignite: TensorFlow Integration





- Use Cases Operational data "High dimension" data (Images, Text, Audio, speech)
- Image data classification
- Natural Language Processing Clinical notes
- Document Classification, Free Form text



To Summarize: Apache Ignite for Continuous Learning at Scale



- Massive Throughput with minimal (or even eliminated) ETL data movements
- Massive operational data sizes + in-place parallel processing
- Faster cycle times for OLTP transactions, ML/DL operations

Integrates with Existing ML / DL operations

- Low-level Distributed APIs to integrate with Auto ML and other Data Science workflows
- Apache Ignite integrations to leverage Spark, TensorFlow pipelines; including Model imports from other tool sets such as XGBoost, Spark, others



Demonstrations, Code Walkthroughs



APIs: Building and Using the Predictive Model



Some representative API call signatures on typical Model object	Pipeline Stage comments
<pre>MyIgniteModel = My<algorithmspecific>Trainer.fit() MyXGBModel = MyModelBuilder.build(, myXGBParser,) MySparkModel = (MyAlgorithm)SparkModelParser.parse()</algorithmspecific></pre>	Create an instance of the Model: MyModel created either by fit() call (distributed training on the algorithm specific trainer) or parsing from Spark, XGBoost, other from saved Ignite Model . Note: Use of the New Pipeline calls means you can replace fit() calls on individual algorithms with a single API if desired
Accuracyvalue = MyEvaluator. evaluate (, MyModel,)	Test the model before you use it: Accuracy value is a number (double). MyModel passed as one of the inputs for accuracy checking. Several accuracy measures available including: accuracy, F-measure, Precision, Recall, Specificity, MAE
<pre>Predictedvalue = MyModel.predict(input field values)</pre>	Predict outputs for incoming transaction: "Predicted value" is a double returned from the deployed model to an external application. You should probably wrap this call in a service to enforce things like security, latency, etc
Measuring Predictions versus Actuals performance has no standard API per se. You'll need to capture both values and then do comparisons periodically	Track Prediction accuracy over time : Examples: For Regression , a common measure for error is MAE (Mean Absolute Error) or MSE (Mean Squared Error). For Classification , (Correct # guesses / total # guesses)
<pre>MyModelVersion2 = MyModelVersion1.update(, newTransactionsCache, MyModelVersion1,)</pre>	When model accuracy drifts over time, update with latest data sample: Input arguments include a reference to the older model plus a reference to the new set of data



Building and Using the Predictive Model: try out these examples, many more....



- app org.apacne.ignite.examples.misc.springbean
- 🖊 🌐 org.apache.ignite.examples.ml
 - D package-info.java
- H org.apache.ignite.examples.ml.tree.randomforest
- 🔠 org.apache.ignite.examples.ml.tutorial
- Harrison of the second seco
- > 🛲 ora.apache.ionite.examples.ml.util

All ML examples contained in this package

Building Models Tutorial package provides best overall overview (Enhanced for 2.8)

- More algorithms added, including hyperparameter tuning steps
- Complete pipeline shows accuracy comparisons for each algorithms: pipeline aggregrator, missing values processor, categoricals processor, scaling, test-train split, cross validation, several trainers...

- > M IrainingwithCustomPreprocessorsExample.java
- - IgniteModelDistributedInferenceExample.java
 - Q IgniteModelDistributedInferenceExample
 - ModelStorageExample.java
 - 🔉 🚺 package-info.java

Using a Model Deploying then invoking

output = mdl.predict(inputrecord)

 Also show loading to reuse a saved Model from storage



28 2019 © GridGain Systems

Python support

 <u>https://github.com/gridgain/ml-python-api</u>

Resources

Documentation:

• Examples and Tutorials:

https://apacheignite.readme.io/docs

- <u>https://github.com/apache/ignite/tree/master/examples/s</u> <u>rc/main/java/org/apache/ignite/examples/ml</u>
- Details on TensorFlow
 - <u>https://medium.com/tensorflow/tensorflow-on-apache-ignite-99f1fc60efeb</u>







Advanced Toolkit for Machine Learning Experts

- Pipelining APIs
 - https://apacheignite.readme.io/docs/pipeline-api
- Evaluators
 - <u>https://apacheignite.readme.io/docs/evaluator</u>
- Model cross-validation
 - https://apacheignite.readme.io/docs/cross-validation
- Models ensembling
 - https://apacheignite.readme.io/docs/ensemble-methods



Join Apache Ignite Community!

- Rapidly Growing Engineering Community
- Great Way to Learn Distributed Systems, Computing, SQL, ML, Transactions
- How To Contribute:
 - https://ignite.apache.org/community/contribute.html
- Join Ignite Meetups:
 - https://ignite.apache.org/meetup-groups.html







Q & A



