

Accelerate AI workloads with Intel® oneDNN & oneDAL

A Hands-on workshop

Description

The Intel® oneAPI Deep Neural Network Library (oneDNN) provides highly optimized implementations of deep learning building blocks. With this open source, cross-platform library, deep learning application and framework developers can use the same API for CPUs, GPUs, or both—it abstracts out instruction sets and other complexities of performance optimization.

Using this library, you can:

- Improve the performance of frameworks you already use, such as Intel® AI Analytics Toolkits, Intel® Distribution for PyTorch*, and Intel® Distribution for TensorFlow*.
- Develop faster deep learning applications and frameworks using optimized building blocks.
- Deploy applications optimized for Intel® CPUs and GPUs without writing any target-specific code.

Intel® oneAPI Data Analytics Library (oneDAL) is a library that helps speed up extensive data analysis by providing highly optimized algorithmic building blocks for all stages of data analytics (preprocessing, transformation, analysis, modelling, validation, and decision-making) in batch, online, and distributed processing modes of computation. The library optimizes data ingestion along with algorithmic computation to increase throughput and scalability. It includes C++ and Java* APIs and connectors to popular data sources such as Spark* and Hadoop*. Python* wrappers for oneDAL are part of Intel Distribution for Python.

The workshop will give a basic understanding of Intel® oneAPI Deep Neural Network Library (oneDNN) and Intel® oneAPI Data Analytics Library (oneDAL) through a quick hands-on session which will provide participants with experience in implementing and optimizing neural networks using these libraries.

We'll see how Popular Deep learning frameworks such as TensorFlow and PyTorch are leveraging Intel® oneDNN and other Intel optimisations i.e. Intel® Extension for PyTorch(IPEX) and Intel® Optimization for TensorFlow. Similarly, for Machine Learning libraries such as Scikit-learn and XGBoost, Intel® oneDAL can be used.

Throughout the workshop, there will be ample time for Q&A and discussion, so that participants can get help with any challenges they encounter and explore the potential applications of oneDNN and oneDAL in their work.

Agenda

- Introduction to Intel® oneDNN and Intel® oneDAL, including a hands-on session to get started on the Intel® oneAPI DevCloud.
- Hands-on exercise to learn how oneDNN optimizations can be leveraged in TensorFlow* and PyTorch* using python APIs and environment variables.
- Hands-on exercise to learn how oneDAL optimizations can be leveraged in Scikit-learn* using python API daal4py.

Key Takeaways

- Participants will gain an understanding of the features and capabilities of these two libraries, and how they can be used to accelerate performance in their applications.
- The hands-on session will give participants an opportunity to get familiar with the Intel® oneAPI DevCloud environment.
- Attendees will learn the basic understanding of the programming model of oneDNN and explore essential concepts and features of oneDAL routines.
- How one can further optimize PyTorch and TensorFlow workloads with Intel® Extension for PyTorch*(IPEX) and Intel® Optimization for TensorFlow.
- The Q&A and open discussion sessions provide an opportunity to ask questions and clarify any doubts.

Pre-Requisites

- A laptop or computer with a chrome/Firefox web browser and a stable internet connection for accessing the workshop materials and the Intel® oneAPI DevCloud environment.
- Basic knowledge of linear algebra.
- Familiarity with programming in Python/C++.
- Familiarity with machine learning and deep learning concepts.

The session will be immensely useful for :

AI Developers, HPC Engineers, AI/ML Enthusiasts, Research scholars, Tech leads, Computer Science Professors, Data Scientists, Programmers, Developers, and Developers who are very new to oneAPI & DevCloud

Session Speaker - Aditya Sirvaiya

AI Software Solutions Engineer @ Intel® Corporation

Aditya Sirvaiya is an AI Software Solutions Engineer in Intel® Artificial Intelligence and Analytics (AIA) group, primarily working on enabling customers and partners to optimize their AI workloads/solutions on Intel platforms using Intel AI software tools and oneAPI high-performance libraries. He is also responsible for supporting and contributing to Intel AI Analytics Toolkits and its components. Aditya holds a bachelor's degree in Engineering

Physics from IIT Delhi and a master's degree in Computer Science with a specialization in AI from IIT Bombay.

LinkedIn: <https://www.linkedin.com/in/aditya-sirvaiya-3390136a/>

Date of Webinar: 19th April 2023

Timings: 5 PM to 7 PM IST

DevCloud Link: <https://devcloud.intel.com/oneapi/>

Discord Link: <https://discord.gg/ycwqTP6>