

# Welcome to Deep Graph Library Tutorials and Documentation

Deep Graph Library (DGL) is a Python package built for easy implementation of graph neural network model family, on top of existing DL frameworks (currently supporting PyTorch, MXNet and TensorFlow). It offers a versatile control of message passing, speed optimization via auto-batching and highly tuned sparse matrix kernels, and multi-GPU/CPU training to scale to graphs of hundreds of millions of nodes and edges.

## Getting Started

For absolute beginners, start with the [Blitz Introduction to DGL](#). It covers the basic concepts of common graph machine learning tasks and a step-by-step on building Graph Neural Networks (GNNs) to solve them.

For acquainted users who wish to learn more advanced usage,

- [Learn DGL by examples](#).
- Read the [User Guide](#) ([中文版链接](#)), which explains the concepts and usage of DGL in much more details.
- Go through the tutorials for [Stochastic Training of GNNs](#), which covers the basic steps for training GNNs on large graphs in mini-batches.
- [Study classical papers](#) on graph machine learning alongside DGL.
- Search for the usage of a specific API in the [API reference manual](#), which organizes all DGL APIs by their namespace.

## Contribution

DGL is free software; you can redistribute it and/or modify it under the terms of the Apache License 2.0. We welcome contributions. Join us on [GitHub](#) and check out our [contribution guidelines](#).

## Index

- [Index](#)