GraphNeuralNetworks.jl: Deep Learning on Graphs with Julia

Carlo Lucibello

CARLO.LUCIBELLO @UNIBOCCONI.IT

Bocconi University Department of Computing Sciences Bocconi Institute for Data Science and Analytics Milan, Italy

Aurora Rossi

Université Côte d'Azur, INRIA, CNRS, I3S Sophia Antipolis, France AURORA.ROSSI@INRIA.FR

Editor: Sebastian Schelter

Abstract

GraphNeuralNetworks.jl is an open-source framework for deep learning on graphs, written in the Julia programming language. It supports multiple GPU backends, generic sparse or dense graph representations, and offers convenient interfaces for manipulating standard, heterogeneous, and temporal graphs with attributes at the node, edge, and graph levels. The framework allows users to define custom graph convolutional layers using gather/scatter message-passing primitives or optimized fused operations. It also includes several popular layers, enabling efficient experimentation with complex deep architectures. The package is available on GitHub: https://github.com/JuliaGraphs/ GraphNeuralNetworks.jl.

Keywords: graph neural networks, Julia, machine learning, deep learning, heterogeneous graphs, temporal graphs

1. Introduction

Graphs are a fundamental data structure for representing complex relationships between entities across diverse domains, including social networks, biological systems, and more (Newman, 2018). Graph Neural Networks (GNNs) are a specialized class of neural networks designed to process graph-structured data, leveraging iterative message passing on the graph topology to capture key patterns and relationships (Scarselli et al., 2008; Kipf and Welling, 2017; Gilmer et al., 2017).

In this paper, we introduce GraphNeuralNetworks.jl, an open-source framework for designing and training graph neural networks. Built for flexibility and ease of use, the package provides extensive functionalities for efficient deep learning on graphs. It supports homogeneous, heterogeneous, and temporal graphs, offering a versatile toolkit for a wide range of applications. The package is implemented in Julia, a high-performance programming language optimized for scientific computing (Bezanson et al., 2017).

 $[\]textcircled{O}2025$ Carlo Lucibello and Aurora Rossi.

Originally launched as a single package, GraphNeuralNetworks.jl has evolved into a suite of four interdependent packages, all hosted within a single GitHub repository. These packages are organized as follows:

- GraphNeuralNetworks.jl: This package provides stateful graph convolutional layers based on the machine learning framework Flux.jl (Innes, 2018). It serves as the primary interface for Flux users and depends on the GNNlib.jl and GNNGraphs.jl.
- GNNLux.jl: This package offers stateless graph convolutional layers designed for the Lux.jl machine learning framework (Pal, 2023). It acts as the main interface for Lux users and relies on GNNlib.jl and GNNGraphs.jl.
- **GNNlib.jl**: This core package implements the foundational message-passing functions and functional graph convolutional layers. It serves as the shared base for building graph neural networks in both the Flux.jl and Lux.jl frameworks. It depends on GNNGraphs.jl.
- **GNNGraphs.jl**: This package provides the graph data structures and utility functions for handling graph data.

From this point forward, we will use GNNs.jl to refer to the overall collection of these packages.

2. Related Graph Neural Network Packages

GNNs.jl draws inspiration from GeometricFlux.jl, the first GNN library for Julia, which is no longer actively maintained at the time of writing. GNNs.jl improves upon it in several key areas, offering a cleaner interface, enhanced performance and flexibility, broader GPU support, and a richer set of features.

In the Python ecosystem, several analogous libraries exist. The most prominent among them, from which GNNs.jl borrows some ideas, are PyTorch Geometric (Fey and Lenssen, 2019), Deep Graph Library (DGL) (Wang et al., 2019), and PyTorch Geometric Temporal (Rozemberczki et al., 2021). While GNNs.jl provides comparable features, it currently lags behind Python alternatives in certain areas, such as large graph training. These limitations are partly due to broader constraints within Julia's deep learning ecosystem, which we anticipate will improve in the near future.

3. Package Design and Implementation

Installation. All packages in the GNNs.jl suite can be installed through the Julia package manager. Flux users should install GraphNeuralNetworks.jl, while Lux users should install GNNLux.jl. Both packages re-export functionality from GNNGraphs.jl and GNNlib.jl, so these core packages do not need to be installed directly.

Graph Data Structures. The primary graph type defined in GNNs.jl is GNNGraph. It encodes the graph topology using either the COO format (default), a sparse adjacency matrix, or a dense adjacency matrix. This type supports features at the node, edge, and

graph levels and can batch multiple graphs together for efficient training on small graph datasets. GNNGraph is fully compatible with the extensive set of graph operations provided by Graphs.jl. Additional graph types include GNNHeteroGraph, which supports graphs with multiple node and edge types, and TemporalSnapshotsGNNGraph, designed for time-varying graphs and dynamic features.

Message Passing. GNNs.jl relies on the message-passing primitives scatter and gather , which are used within the functions apply_edges and reduce_neighborhood. With KernelAbstractions.jl, a single high-level kernel for each of these primitives can be written in Julia and JIT-compiled for multiple backends. This enables users to define custom layers using the message-passing framework (Gilmer et al., 2017):

```
m = apply_edges(message, g, xi, xj, eji) # apply message function on each edge

m = aggregate_neighbors(g, op, m) # reduce operation on each neighborhood
```

Moreover, the propagate function combines these two operations, fusing them for efficiency when possible. Leveraging Julia's multiple dispatch, fused versions of the gather/scatter operations (generalized matrix multiplication) can be implemented for specific message functions and aggregation operations, further optimizing performance.

Graph Convolutional Layers. GraphNeuralNetworks.jl offers a variety of popular graph convolutional layers, including GAT (Veličković et al., 2017) and GIN (Xu et al., 2019), among others¹. In addition to these built-in layers, users can define custom layers and combine them into complex models, as demonstrated in the code below.

```
struct GraphConv <: GNNLayer</pre>
   weight1
    weight2
    bias
    \sigma
    aggr
end
Flux.@layer GraphConv
function GraphConv(ch, \sigma=identity; aggr=+)
    din, dout = ch
    W1 = glorot_uniform(dout, din)
    W2 = glorot_uniform(dout, din)
    b = zeros(Float32, dout)
    return GraphConv(W1, W2, b, \sigma, aggr)
end
function (l::GraphConv) (g::GNNGraph, x)
   m = propagate(copy_xj, g, l.aggr, xj=x)
    x = 1.weight1 * x .+ 1.weight2 * m
    return 1.\sigma.(x . + 1.bias)
end
```

```
struct Model
   conv1::GraphConv
    conv2::GraphConv
    dense::Dense
end
Flux.@layer Model
function Model(din, dout)
   Model(GraphConv(din => 64, relu),
          GraphConv(64 => 64, relu),
          Dense(64 => dout))
end
function (m::Model)(g::GNNGraph, x)
   h = relu.(m.convl(g, g.x))
   h = relu.(m.conv2(h))
   return m.dense(h)
end
```

^{1.} A complete list of available layers can be found at https://juliagraphs.org/ GraphNeuralNetworks.jl/graphneuralnetworks/api/conv for the Flux.jl backend and https://juliagraphs.org/GraphNeuralNetworks.jl/gnnlux/api/conv for the Lux.jl backend.

Batch Processing, GPU Support, and Datasets. The package provides batch processing capabilities for efficiently handling large datasets and supports GPU acceleration (CUDA and AMDGPU at the time of writing) by leveraging the JuliaGPU ecosystem to significantly speed up computations (Besard et al., 2018, 2019). Additionally, the package integrates seamlessly with real-world graph datasets available through MLDatasets.jl.

Training Example. The following code demonstrates how to build a simple GNN for a regression task using a synthetic dataset. This illustrative example highlights the definition of the model, setup of the optimizer, creation of data loaders, and training of the model over multiple epochs. It also demonstrates the use of batching and GPU acceleration with CUDA for enhanced performance.

```
using GraphNeuralNetworks, Flux, Statistics, MLUtils, CUDA
make_graph() = rand_graph(n, m,
   ndata = (; x = randn(Float32, 16, n)), gdata=(; y = randn(Float32)))
train_data = [make_graph() for i=1:num_graphs]
train_loader = DataLoader(train_data, batchsize=32, shuffle=true, collate=true)
model = GNNChain(GCNConv(16 => 64),
                 BatchNorm(64), # normalize across node dimension
                 x \rightarrow relu.(x),
                 GCNConv(64 \Rightarrow 64, relu),
                 GlobalPool(mean), # Aggregate for graph-wise prediction
                 Dense(64, 1)) |> gpu
opt_state = Flux.setup(Adam(1f-4), model)
loss(model, g) = mean((vec(model(g, g.x)) - g.y).^2)
for epoch in 1:num_epochs
    for g in train_loader
        g = g \mid > gpu \# transfer the batch of graphs to gpu
        grad = gradient(model -> loss(model, g), model)
        Flux.update!(opt_state, model, grad[1])
    end
end
```

4. Conclusions and Future Plans

Leveraging the unique features of the Julia language, effectively solving the two-language problem, and the rapidly evolving deep learning and GPU programming Julia ecosystems, GNNs.jl has quickly established itself as a versatile and comprehensive suite of packages for building and training graph neural networks in Julia. It supports a wide variety of graph types, including heterogeneous and dynamical ones, and provides a robust set of graph convolutional layers that integrate seamlessly with the Lux.jl and Flux.jl deep learning frameworks. Future plans for the package include extending support to additional hardware backends, such as Apple Silicon Metal and TPUs, incorporating XLA compilation via Reactant.jl, and enabling multi-GPU training. Further development is also planned to optimize message-passing operations and improve support for sparse matrix operations on GPUs.

Acknowledgments

A.R. has been supported by the French government, through the France 2030 investment plan managed by the Agence Nationale de la Recherche, as part of the "UCA DS4H" project, reference ANR-17-EURE-0004, NUMFocus organization and Google Summer of Code program 2023. C.L. acknowledges funding from the European Union – Next Generation EU (PRIN 2022 project 202234LKBW "Land(e)scapes" and PRIN PNRR 2022 project P20229PBZR "When deep learning is not enough".)

References

- Tim Besard, Christophe Foket, and Bjorn De Sutter. Effective extensible programming: Unleashing Julia on GPUs. *IEEE Transactions on Parallel and Distributed Systems*, 2018. ISSN 1045-9219. doi: 10.1109/TPDS.2018.2872064.
- Tim Besard, Valentin Churavy, Alan Edelman, and Bjorn De Sutter. Rapid software prototyping for heterogeneous and distributed platforms. Advances in Engineering Software, 132:29-46, 2019. ISSN 0965-9978. doi: https://doi.org/10.1016/j.advengsoft. 2019.02.002. URL https://www.sciencedirect.com/science/article/pii/ S0965997818310123.
- Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. SIAM review, 59(1):65–98, 2017. URL https://doi.org/ 10.1137/141000671.
- Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In ICLR 2019 Workshop on Representation Learning on Graphs and Manifolds, 2019. URL https://arxiv.org/abs/1903.02428.
- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 1263–1272. JMLR.org, 2017.
- Mike Innes. Flux: Elegant machine learning with julia. *Journal of Open Source Software*, 2018. doi: 10.21105/joss.00602.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=SJU4ayYgl.
- Mark Newman. Networks. Oxford university press, 2018.
- Avik Pal. Lux: Explicit Parameterization of Deep Neural Networks in Julia, April 2023. URL https://doi.org/10.5281/zenodo.7808904.
- Benedek Rozemberczki, Paul Scherer, Yixuan He, George Panagopoulos, Alexander Riedel, Maria Astefanoaei, Oliver Kiss, Ferenc Beres, Guzman Lopez, Nicolas Collignon, and Rik

Sarkar. PyTorch Geometric Temporal: Spatiotemporal Signal Processing with Neural Machine Learning Models. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, page 4564–4573, 2021.

- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1): 61–80, 2008.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv preprint arXiv:1909.01315, 2019.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ryGs6iA5Km.