
Just-in-Time Dynamic-Batching

Sheng Zha¹, Ziheng Jiang², Haibin Lin¹, Zhi Zhang¹

¹Amazon AI, Amazon Web Services

²Paul G. Allen School, University of Washington

{zhasheng, haibilin, zhiz}@amazon.com

ziheng@cs.washington.edu

Abstract

Batching is an essential technique to improve computation efficiency in deep learning frameworks. While batch processing for models with static feed-forward computation graphs is straightforward to implement, batching for dynamic computation graphs such as syntax trees or social network graphs is challenging due to variable computation graph structure across samples. Through simulation and analysis of a Tree-LSTM model, we show the key trade-off between graph analysis time and batching effectiveness in dynamic batching. Based on this finding, we propose a dynamic batching method as an extension to MXNet Gluon’s just-in-time compilation (JIT) framework. We show empirically that our method yields up to 6.25 times speed-up on a common dynamic workload, a tree-LSTM model for the semantic relatedness task.

1 Introduction

As modern neural networks evolve, dynamic neural network structures such as variable-length sequences [5, 14], trees [16, 15], and graphs [7, 4] become increasingly important. For instance, social networks have complex graph structures, and natural language processing (NLP) problems often have variable-length sequences, each accompanied by a different parse tree [16]. While most deep learning frameworks (e.g. Tensorflow [1], PyTorch [13], dyNet [11], and Apache MXNet [3]) can express these dynamic neural networks as computation graphs, training with dynamic graphs can be much slower compared to static ones. This is because for each sample, the framework needs to construct different computation graphs, move data across the memory hierarchy, which limits parallelism and efficient utilization of hardware resources. Also, optimization techniques such as memory planning might not be effectively reused due to the variations in the computation graphs. Batch processing amortizes these overhead, provides better data locality, and enjoys better parallelism. For example, many vector-vector multiplication operations that share the first operand can be batched into one matrix-vector multiplication.

However, naive batching technique is often accompanied by additional data preprocessing steps such as padding and bucketing for mini-batching variable-length sentences. Models with dynamic graph structures often require more complex data preprocessing, such as ingesting parse tree data according for a tree-based model. In this paper, we focus on dynamic batching methods that can automatically rewrite computation graphs to enable batching without extra data preprocessing, in a just-in-time fashion.

Given a computation graph, the goal of dynamic batching is to identify subgraphs that can be executed together. These subgraphs must be isomorphic in the sense of same graph topology, same node types, and same parameterization. The isomorphism check guarantees consistent results between the original computation graph and our proposed batching method. Each subgraph must also be capable of performing data-parallel (i.e. single instruction, multiple threads, SIMT) computation for a stack of samples to enable rewriting them into a single batched subgraph.




Figure 1: Example computation graphs for Tree-LSTM. The outer frame marks the subgraph (tree) boundary. Each node represents an operator, and each *op* node represents the same operation with the same parameterization. When analyzing these graphs at operator level, all leaf nodes can be batched. However, when analyzed at the subgraph level, C2 and C3 cannot be batched because they are not isomorphic.

An unfortunate fact is that the analysis to identify isomorphic nodes at fine-grain require significantly more time than coarse-grain analysis. For example, consider the computation graph in Figure 1. In this case, all leaf nodes are of the same type and can be batched together. Batching all leaf nodes results in the maximum computation reduction, but the isomorphism check is required on all 7 nodes. On the other hand, if we take the perspective of a lower granularity and consider the two subgraphs C2 and C3, we only need to compare the two subgraphs with sum nodes as roots. This reduces the analysis overhead but C2 and C3 are no longer co-exist in same batch. Therefore, selecting the right granularity is vital to the efficiency of dynamic batching algorithms.

In this paper, we experimentally analyze the trade-off between graph processing time and individual subgraph batching speed-up. Through simulation we show that on a Tree-LSTM [16] model, it is important to utilize the available granularities in user code and its effectiveness is obvious. Based on this insight, we present just-in-time batching, a dynamic batching method well-suited for the just-in-time compilation in Apache MXNet Gluon [3]. Our experiments indicate that our method performs significantly better than traditional batching methods by fully utilizing the subgraph structures under guidance in user code, with only one line of code change.

2 Related Work

Some previous work approaches the batching problem from an application-specific perspective, such as cellular batching for RNNs with variable length inputs [6], and batched graph execution [17]. There also exist several general purpose dynamic batching methods, including Tensorflow Fold [8], DyNet [12], MatchBox [2], and Cavs [18]. In this work, we focus on the general purpose methods.

Tensorflow Fold rewrites the computation graph before execution by matching graph node depth, and converts dynamic graphs into graphs with control-flow operations. This approach is widely applicable but limited in two ways. Due to the granularity chosen at the time of graph constructions (i.e. subgraph level batching), this approach cannot batch operators that exist at finer granularity (e.g. addition and matmul). As a result, some subgraphs cannot be batched even if they only vary in minor ways, such as trees with variable number of children (Figure 1). Fold would treat such variants as completely independent components in the graph and produce inefficient rewritten graphs. Furthermore, because the graph transformation happens before execution, this approach is less applicable when workload appears incrementally at irregular cadence while previous load is still being executed. Such workload is commonly seen in model serving, where model inference requests can appear at any time. By performing dynamic batching as part of JIT, our approach can handle such cases with good batching efficiency.

DyNet auto-batching approach overcomes the above limitations by analyzing at the operator level online during execution. The auto-batching analysis is undertaken in the scheduler. Kernels with the same signatures within the frontier of the running operators will be batched if they satisfy input dependencies. In order to make sure that there are enough nodes in the pool to construct a potentially large batch, this approach depends on a heuristic to decide whether to wait for more nodes to arrive at the frontier before executing. For workloads that involve many kernels, the analysis overhead can become a bottleneck and cannot be hidden through asynchronous execution, and thus dominates the overall processing time. In contrast, benefit from various granularities available in MXNet Gluon, our approach enjoys freedom to choose a lower granularity to reduce the run-time analysis overhead.

Table 1: Statistics of kernels and subgraphs for Tree-LSTM on Sentences Involving Compositional Knowledge (SICK) dataset. The nodes in the trees from SICK dataset have varying number of children between 0 and 9. **No-batch**: the kernel launch count if no batching happens. **Batch**: the kernel launch count if batching happens. **Ratio**: batching ratio. The batched count for kernels is calculated from input graphs in the dataset. Count for subgraph batching is observed through simulation.

Nodes	kernel	subgraph
No-batch	5018658	148681
Batch	~2650	1081
Ratio	1930x	137x

MatchBox takes a similar approach to Tensorflow Fold but in the context of imperative execution in PyTorch. At the time when MatchBox was devised, PyTorch executes the operators in a blocking way. Thus, MatchBox creates stubs for all functional calls to collect the AST for batching analysis. During execution, the computation graphs are dynamically constructed at runtime. Our approach can instead happen as part of the Gluon JIT and cache the rewriting of graphs.

Cavs [18] proposes a vertex-centric programming interface that decompose a dynamic neural network with two components: a static vertex function that can be batched together, and a sample specific graph which is dynamic. Our approach, built on top of MXNet Gluon’s JIT, does not force models to be programmed as vertex and graph functions, and only requires one line of code change to enable batching.

3 Effects of Subgraph Granularity on Dynamic Batching

To demonstrate the significance of the choice on granularity, we use a Tree-LSTM model in a simulation and compare the effectiveness and analysis overhead from two different choices of granularity: kernel level and subgraph (cell) level. In the simulation, we apply the Fold method to batch 256 samples at a time, and record the effective batching-ratio, as well as the zoom-ratio by varying the granularity. With same workload, batching ratio is defined as the ratio of the number of kernel launches with and without batching techniques. The larger batching ratios achieved, more operators or subgraphs are being executed together, therefore higher speed-up can be achieved.

As shown in Table 2, we observe orders of magnitude difference between no-batch counts of kernels and subgraph, respectively. The difference between the batch ratios (1930x vs 137x) of kernels and subgraphs is strong evidence that the benefit of batching at fine-grain can be significant. Therefore, it is vital to pick a granularity to balance reduced the analysis overhead and potential batching speed-up benefits.

In order to understand such difference, we closely examine the Tree-LSTM structure. Each LSTM cell variant consists of 33 operators, where only 4 operators would vary based on the number of children. The rest of the operators have the exact same graph structure. Due to the fact that cells with different number of children cannot be batched, the batching capabilities of the rest operators are ruined by these four dynamic operators.

4 Just-in-time Dynamic Batching

There are several factors to consider in the design of the dynamic batching methods. From the analysis on Tree-LSTM model (Section 3), we recognize that the key to high performance batching scheme is choosing the right granularity. Such granularity are often readily available from user code, because programmers are reasonably good at organizing code in a modular way to reduce repetition. Gluon HybridBlock supports user-defined subgraphs at various levels, therefore we can take advantage of it to decide batching granularity during analysis (Section 4.1). Our batching method also supports imperative execution in Gluon, which is integrated into Gluon’s JIT 4.2.

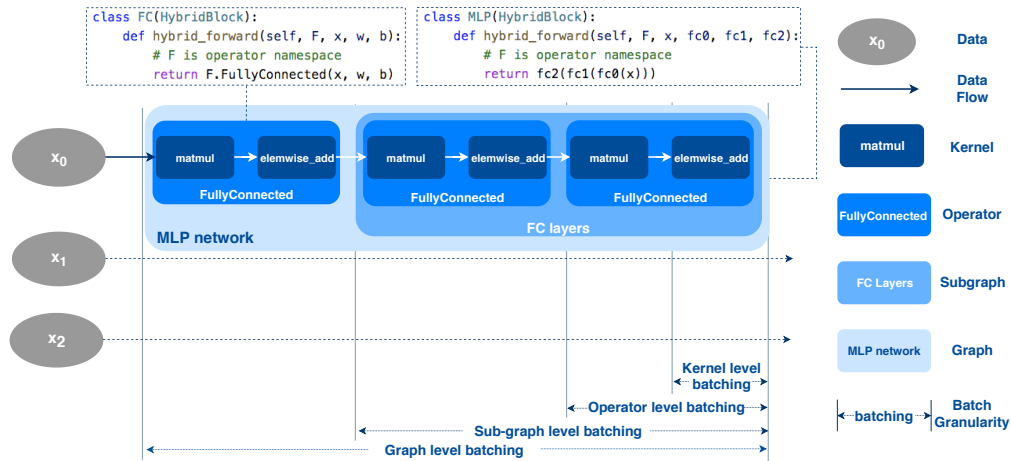


Figure 2: Illustration of batching granularity in a multi-layer perceptron network. Individual data is passed through stacked fully connected layers (activation layers are omitted for simplicity). Gluon supports multiple granularity separation levels. Aside from graph-level batching which is traditional batching mechanism, subgraph, operator, and kernel level batching provide more granularity potentials for optimization.

4.1 Analysis Granularity

We visit the dynamic batching problem with MXNet’s Gluon interface as it provides great potentials in analyzing and fulfilling different levels of batching granularity. The Gluon interface offers a useful high-level abstraction, called "HybridBlock", which can be used as building block for the repeated computation subgraph for models with dynamic computation graphs, such as RNN cells. Through just-in-time compilation (JIT), it creates and caches the recurrent cell logic into a symbolic graph, which is passed to execution engine as an operator node with a subgraph inside.

The subgraph is a loosely tied operator union (e.g. two layers of consecutive fully-connected layers) that can be cache-and-batched as a whole processing unit. The operators again usually consist of series of instructions that can be further divided into finer-grained computations. If needed, a fully connected operator can be further divided into two kernels: matrix multiplication and element-wise addition. Such flexibility not only enables the versatile needs of expressing models, but also allows us to pick the level of granularity we need.

4.2 Lazy Execution for Imperative Mode

In addition, Gluon supports imperative execution, on top of which JIT converts to symbolic execution mode, and optimize and cache the computation graph. Thus, our batching method should also support both deferred imperative execution, and operate in symbolic mode. Supporting dynamic-batching in the symbolic programming style is straightforward, because it is possible to examine the complete computation graphs for all samples. However, this does not hold for supporting imperative execution.

The NDAarray interface in MXNet supports eager execution, in which we can only examine one small part of the computation graph at a time. We extend NDAarray to support lazy execution through a new interface called NDAarrayFuture. Unlike regular NDAarray which schedules the computation as soon as requests arrive, this interface allows delaying computation until user code requests to collect any of the resulting arrays. This interface benefits from the great usability from regular imperative programming, such as easy debugging since users can request for the values of any array at anytime. By delaying computation, it also sees more complete computation graphs, thus leading to more opportunity to optimize.

The operator registration mechanism in MXNet allows us to easily get an overview of all supported operators, and generate stub code according these operators, so that these NDAarrayFuture instances can behave exactly like regular NDAarrays but can be evaluated lazily. Each time a new NDAarrayFuture reference is created through the generated interface, we save the corresponding computation and

organize the nodes and the input arguments in a look-up table according to their depth. The nodes at the same depth are independent of each other and thus can be evaluated in parallel. In order to identify the nodes that can be batched together, we use the computation node type, the node settings, the input argument layouts, as well as result look-up index to form a unique look-up key.

Further, we devise a new dynamic batching scope to allow users to easily specify the computation and the number of samples that we should try to batch. When entering such scope, the interface automatically create the aforementioned look-up table, whose references are kept in the `NDArrayFuture`. Within the scope, the execution of the computation of all the `NDArrayFuture` is delayed, until the code reaches the exit of the scope. When exiting the scope, the computation is started, and is evaluated in the order of the computation graph depth. The execution results are actual `NDArrays`, and are filled in the look-up table. After exiting the batching scope, the `NDArrayFuture` can refer to the actual `NDArray` results and behave like regular `NDArrays`.

4.3 Dynamic Batching

Once we collected the graphs, we reorganize them into a look-up table so that the computation nodes that can be batched together reside in the same slot. When executing, we stack these samples on the batch axis, and feed the newly formed batch to the common subgraph and launch execution. After execution, we slice the output `NDArray` to obtain the results that correspond to individual samples. Since such graph-rewrite can be expressed with symbolic execution, and `MXNet` supports control flows as operators, the graph rewriting can be cached and stored for next forward pass. This also means that through delayed execution, we make dynamic batching as part of the JIT optimization in `MXNet`. Last but not least, `MXNet` currently expresses control-flow also with operators, which means they are also nodes in the computation graphs. Control-flow operations such as for-loop are often used to express repetition, exposing its inner details to the analysis logic is also important to achieve the best performance.

In our implementation, the user only need to add one extra line of code to declare the batching scope to enable dynamic batching, as shown in the following pseudo python code block.

```
1 #network and loss definition
2 net = GraphConvolutionNet()
3 loss = SoftmaxCELoss()
4 #batching scope
5 with mx.batching():
6     #iterate through samples in the batch
7     for data, label in data_batch:
8         #forward computation
9         out = net(data)
10        ls = loss(out, label)
11        #backward computation
12        ls.backward()
13 #parameter update
14 trainer.step()
```

5 Experiments

In this section, we benchmark the training and inference speed of tree-structured LSTM networks on the semantic-relatedness task on SICK dataset [10] for both training and inference. SICK dataset consists of 4500 pairs of sentences, each labeled with a semantic-relatedness score for the two sentences. We use Stanford Parser [9] for extracting the parse trees of the sentences. The nodes in the parse trees from SICK dataset have varying number of children between 0 and 9.

The benchmark runs on EC2 c4.8xlarge instance with Intel Xeon E5-2666 v3 (Haswell) processors. The per-instance method evaluates one input instance at a time, while JIT dynamic batching uses a batch size of 256. With dynamic batching we are able to achieve 5.96x speedup for training and 6.25x speedup for inference.

Table 2: Training and inference speed of Tree LSTM on SICK dataset.

Method	Training (samples/s)	Inference (samples/s)
Per instance	33.77	50.46
JIT dynamic-batching	201.11 (5.96x)	315.54 (6.25x)

6 Conclusion

Deep learning models with dynamic computation graphs are increasingly important in the research community. In this paper, we identify the key design choice for dynamic batching algorithms, the granularity of the subgraph for analysis, which affects both analysis time and discover-ability of computation available for batching. Based on the design of JIT in MXNet, we design a batching method that is suitable for both imperative execution and symbolic execution under JIT. Our method utilizes the existing hierarchy in user code, which provides good starting point in terms of subgraph granularity. Our experiments show the effectiveness of this simplistic approach.

Furthermore, the asynchronous extension of NDAarray, i.e. our new NDAarrayFuture interface, provides an extensible basis for other computation optimization such as operator fusion, function approximation. Its flexibility will facilitate the prototyping of optimization techniques in future work.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016.
- [2] James Bradbury. Matchbox. 2018. URL <https://github.com/salesforce/matchbox/>.
- [3] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems. 12 2015. URL <http://arxiv.org/abs/1512.01274>.
- [4] Hanjun Dai, Zornitsa Kozareva, Bo Dai, Alex Smola, and Le Song. Learning steady-states of iterative algorithms over graphs. In *International Conference on Machine Learning*, pages 1114–1122, 2018.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018. URL <http://arxiv.org/abs/1810.04805>.
- [6] Pin Gao, Lingfan Yu, Yongwei Wu, and Jinyang Li. Low latency rnn inference with cellular batching. In *Proceedings of the Thirteenth EuroSys Conference*, page 31. ACM, 2018.
- [7] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [8] Moshe Looks, Marcello Herreshoff, DeLesley Hutchins, and Peter Norvig. Deep Learning with Dynamic Computation Graphs. *Neurosurgery*, 60(5):815–824, 2 2017. ISSN 0148396X. doi: 10.1227/01.NEU.0000255452.20602.C9. URL <http://arxiv.org/abs/1702.02181>.
- [9] Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60, 2014. URL <http://www.aclweb.org/anthology/P/P14/P14-5010>.
- [10] Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. A sick cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and*

- Evaluation (LREC-2014)*. European Language Resources Association (ELRA), 2014. URL <http://www.aclweb.org/anthology/L14-1314>.
- [11] Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, Kevin Duh, Manaal Faruqi, Cynthia Gan, Dan Garrette, Yangfeng Ji, Lingpeng Kong, Adhiguna Kuncoro, Gaurav Kumar, Chaitanya Malaviya, Paul Michel, Yusuke Oda, Matthew Richardson, Naomi Saphra, Swabha Swayamdipta, and Pengcheng Yin. DyNet: The Dynamic Neural Network Toolkit. 1 2017. URL <http://arxiv.org/abs/1701.03980>.
 - [12] Graham Neubig, Yoav Goldberg, and Chris Dyer. On-the-fly Operation Batching in Dynamic Computation Graphs. (Nips):1–13, 2017. URL <http://arxiv.org/abs/1705.07860>.
 - [13] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
 - [14] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018.
 - [15] Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.
 - [16] Kai Sheng Tai, Richard Socher, and Christopher D. Manning. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1556–1566, 2015. ISSN 9781941643723. doi: 10.1515/popets-2015-0023. URL <http://arxiv.org/abs/1503.0075>.
 - [17] Minjie Wang, Lingfan Yu, Quan Gan, Da Zheng, Yu Gai, Zihao Ye, Mufei Li, Jinjing Zhou, Qi Huang, Junbo Zhao, Haibin Lin, Chao Ma, Damon Deng, Qipeng Guo, Hao Zhang, Jinyang Li, Alexander J Smola, and Zheng Zhang. Deep graph library, 2018. URL <http://dgl.ai>.
 - [18] Hao Zhang, Shizhen Xu, Graham Neubig, Wei Dai, Qirong Ho, Guangwen Yang, and Eric P Xing. Cavs: A vertex-centric programming interface for dynamic neural networks. *arXiv preprint arXiv:1712.04048*, 2017.