

---

# Automatic Discovery of Composite SPMD Partitioning Strategies in PartIR

---

Sami Alabed \*    Dominik Grewe    Juliana Franco    Bart Chrzaszcz

Tom Natan    Tamara Norman    Norman A. Rink

Dimitrios Vytiniotis    Michael Schaarschmidt

DeepMind

## Abstract

Large neural network models are commonly trained through a combination of advanced parallelism strategies in a single program, multiple data (SPMD) paradigm. For example, training large transformer models requires combining data, model, and pipeline partitioning; and optimizer sharding techniques. However, identifying efficient combinations for many model architectures and accelerator systems requires significant manual analysis. This work presents an automatic partitioner that identifies these combinations through a goal-oriented search. Our key findings are that a Monte Carlo Tree Search-based partitioner leveraging partition-specific compiler analysis directly into the search and guided goals matches expert-level strategies for various models.

## 1 Introduction

The rapid rise of large neural networks with significant memory and compute requirements have made partitioning strategies critical for enabling their training and inference. For example, large language models such as GPT-3 [2], Gopher [22], PaLM [5] or Chinchilla [10] have relied on various combinations of Megatron-style layer sharding [28, 20], batch/data parallelism, pipeline parallelism [19, 35, 11], or ZeRO sharding/offloading [23, 25].

Model sizes across applications have been outpacing device memory and FLOPS growth: Google TPU [14, 15] v2 reports 46 TFLOPS / 16 GB of HBM per chip, while v4 reports 275 TFLOPS / 32 GB<sup>2</sup> over a period where model parameters – and consequently FLOPS requirements – increased by approximately 4 orders of magnitude. The use of sophisticated parallelism strategies is, therefore, unavoidable. Despite progress in the development of partitioning APIs (e.g., GSPMD-style [34, 18] partitioning annotations exposed in front-end APIs by JAX [1]) or parallelism frameworks such as DeepSpeed [24], customizing new research models to hardware configurations often still requires expert analysis to account for model-specific compute and communication patterns. Recent work has proposed automated parallelism frameworks specifically targeting large neural network training. For example, Alpa [37] defines a hierarchical space of intra- and inter-operator parallelism and solves sub-problems through integer linear programming (ILP). Unity [29] as an extension of FlexFlow [13] and TASO [12] jointly optimizes parallelism and algebraic transformations through hierarchical search. There exists work on applying RL for the partitioning problem [31], recently combining RL with a constraint solver and achieving good generalization properties [32].

---

\*Correspondence: sa894@cam.ac.uk. Work done during an internship at DeepMind.

<sup>2</sup><https://cloud.google.com/tpu/docs/system-architecture-tpu-vm>

We approach automated parallelism with a different philosophy. Our automated SPMD partitioner is closely integrated with our PartIR compiler stack. PartIR [30] provides our partitioner with the ability to validate partitioning decisions, expose a host of static analyses about the module, automatically propagate sharding decisions across a module, and produce analytical cost model estimates. The latter is obtained using a simulator of a lowered and communication-optimized partitioned module at the compiler level. By contrast, ILP-based encodings [37] attempt to address scalability issues in large graphs by e.g. heuristically grouping operators to prune the search space effectively – furthermore, there may be communication optimizations post-ILP that cannot be captured when deciding the parallel plan. FlexFlow [13] addresses scalability by making the simulator in their MCMC search incremental and relies on pre-defined layers/block-operators (e.g. Multi-Head Attention). On the other hand, our approach can partition *any* array operation of a low-level tensor IR (XLA) that high-level NN libraries compile. As we show, our tight compiler integration allows our search to be data and size efficient and run routinely as part of our user (typically researchers) iterative workflows.

In this work, we extend our automated partitioner, Automap [27], to support partitioning a model across multi-dimensional device topology and discover expert-level composite SPMD partitioning strategies. Our contribution is a design of a goal-oriented Monte Carlo tree search (MCTS) [3] that decomposes the problem into smaller sub-problems. Furthermore, we incorporate a partitioning-specific compiler analysis into the MCTS to reduce both the nodes and edges of the tree, improving the search’s robustness. We show that our partitioner discovers composite expert-level SPMD strategies for common models, such as Transformers and GraphNets. Moreover, it produces significantly better than human-written strategies for models as UNet, for which no known strategy is available.

## 2 Background

### 2.1 Logical Meshes and Composite Strategies




Figure 1: The composite strategy of batch and model parallelism over mesh  $\{\text{batch}:N, \text{model}:M\}$ . We also show communication patterns that may emerge; on the left, possible communication along the model axis (e.g., Megatron activation reductions), and on the right, communication along the batch axis (e.g., gradient reductions). The color coding denotes the unique parameter shards that each device holds; e.g., all devices along the batch axis holds the same shard of parameters.

Leveraging composite partitioning techniques has enabled the training of recent large models [2, 22, 10]. The main idea is to structure the available accelerator devices into an n-dimensional logical *mesh* (which will typically, but not necessarily, correspond to the physical communication topology). For instance, we may view 32 TPU devices as a 4x8 mesh or a system of 2 GPU servers with 16 GPUs each as a 2x16 mesh. Once we are given such a logical mesh of available devices, e.g., with a 2D mesh, a conventional strategy would be to do batch parallelism over one axis and parameter sharding (model parallelism) over the other. Figure 1 graphically depicts this strategy. ZeRO-style sharding of the optimizer [23] (on top of batch parallelism, and possibly also parameter sharding) is simply a different stage that shards the *optimizer parameters* along the axis used for *batch parallelism*.

Conceptually, each stage of a composite strategy, like the above, optimizes for specific objectives. For example, batch parallelism and parameter sharding typically improve runtime (while also reducing memory requirements); whereas ZeRO-style optimizer sharding aims mainly towards improving memory (but may improve runtime too, as the typically memory-bound vector operations of the optimizer update are sharded, as was already observed in precursor work [33]). ZeRO “stage-2” optimizer sharding may not increase communication cost since it replaces AllReduce operations with pairs of parameter AllGather and gradient ReduceScatter. ZeRO “stage-3” aims to further improve memory by introducing a separate parameter AllGather for the forward and the backward pass of a model; hence may slightly increase the runtime in favor of keeping smaller tensors live across forward and backward computation. Note that, in our setting, the logical mesh will be given ahead of time by the user; the partitioner’s task is to discover composite strategies like those described above based on user-provided objectives.

```

// worklist = {%arg0, %arg1, %arg2}
func f(%arg0, %arg1, %arg2) {
  %0 = add(%arg0, %arg1)
  ...
}

// worklist = {%arg2}
func f(%arg0, %arg1, %arg2) {
  %0 = tile 0 "x" (%r : range<N>) {
    %slice0 = slice 0 %arg0[%r]
    %slice1 = slice 0 %arg1[%r]
    %a = add(%slice0, %slice1)
    yield %a
  }
  ...
}

```

Figure 2: Sharding argument `%arg0` along axis `batch` in the left module causes `%arg1` to also become sharded; leaving us with a worklist of just `%arg2` and the rewritten module on the right.

## 2.2 PartIR rewriting for SPMD partitioning

The PartIR compiler infrastructure accepts a user-defined ML model in the form of an MLIR [17] module, a user-provided mesh, and exposes rewrite actions on that module to express forms of shardings. The key idea behind PartIR rewriting is to express parallelism of array operations, like `z = matmul(x, w)`, as parallel *tiling* or *reduction* loops, specifically tagged with the axis over which these parallel loops range. e.g:

```

// Tiling loop derived from sharding x
// along dimension 0 on "batch" axis.
%z = tile 0 "batch" (%r : range<N>) {
  %xs = slice 0 %x[%r]
  %zs = matmul(%xs, %w)
  yield %zs
}

// Reduction loop derived from sharding
// both inputs along a contracting
// dimension on "model" axis.
%z = sum "model" (%r : range<M>) {
  %xs = slice 1 %x[%r]
  %ws = slice 0 %w[%r]
  %zs = matmul(%xs, %ws)
  yield %zs
}

```

Compiler-exposed sharding actions, of the form `partition(parameter_id, dimension, axis)`, introduce and *propagate* such tiling loops across the MLIR module, based on powerful propagation rules that depend on the semantics of the array operations in hand.

**Search with PartIR actions** Automap [27] applies such actions to the module, guided by a *worklist* consisting of the module arguments (e.g., model parameters, optimizer state and the data tensors passed to a training step function). Suppose the compiler detects that the sharding of an argument can be propagated to another argument. In that case, both are removed from the worklist and no longer considered for further sharding actions (See Figure 2 for an example.)

As a result, a very small number of sharding actions can lead to high-quality partitions (e.g., low communication overhead) within a reasonable search budget. This can be amplified by further compressing the worklist with user-provided groupings of parameters that should be equi-sharded (e.g. all transformer blocks across a large transformer model). The overall search is driven by a cost model, typically simulated runtime cost and peak per-device memory obtained statically.

## 3 Multi-axis search

In this work, we extend Automap [27] to support multi-axis automatic partitioning, a capability that our PartIR stack already supports through *nesting* tiling and reduction loops, as well as multi-axis loop propagation. However, multi-axis partitioning increases the search complexity exponentially for every axis. The search must consider the dependency between partitioning decisions and the argument for partition, axis, and dimension. Compiler static analysis together with Automap’s [27] worklist mechanism described above helps us prune the partitioning decision branching factor. This section outlines the major additions needed to reduce the complexity of the problem and reach fast expert-level sharding on several representative models.

### 3.1 Monte Carlo Tree Search for SPMD compilation

The automated partitioner is implemented as a Monte Carlo Tree Search (MCTS) [3] with upper confidence bound for trees (UCT). The MCTS starts with no prior knowledge about the model’s

partition, and by exchanging with PartIR partitioning actions for performance estimates, the MCTS builds a policy that allows it to explore the partitioning space efficiently. Each action in the MCTS translates to a one-hot encoding of a partitioning decision (argument, dimension and mesh axis). The reward is the estimated performance calculated by PartIR and observed at the end of an episode.

The main users of our partitioner are researchers who partition their models during iterative research workflows. In the next section, we discuss how the search can discover effective solutions quickly to enable fast research iterations.

### 3.1.1 Domain specific enhancements over Automap

As in Automap [27], we leverage user equi-sharding annotations to partition repeated layers all at once, as if partitioning a single layer, reducing the search space further. Moreover, we introduce several domain-specific enhancements into the MCTS to improve its efficiency and reduce the complexity of the problem in preparation for the multi-axis search.

**Compressed state representation** In the original Automap work, [27], the sequences of partitioning decisions identified the module’s state in the MCTS. However, PartIR infers and propagates partitioning decisions of a single argument to other values across the module. For example, in Figure 2, we can arrive to the module and worklist on the right by *either* partitioning %arg0 or %arg1. This leads to an interesting observation; there could be several permutations of decisions within the same axis that lead to a similar final partitioned model and increase the branching factor of the search needlessly. PartIR provides a static analysis of the state of each variable of the model (e.g., on which axes and dimensions it has already been sharded). We leverage this analysis to identify how partitioning actions impact the module’s worklist. Such impact is then used to identify the state of the module. This enhancement compresses the state-space of the MCTS, as there are fewer *truly* unique states - akin to reducing the number of nodes in the tree, allowing the search to explore previously infeasible regions to reach.

**Action grouping** Building on the previous finding, various partitioning decisions (actions) can lead to the same state after propagation. This knowledge would save the search from having to take action and find that out - akin to collapsing the out edge from the node, and effectively further reducing the branching factor of the tree [4]. To achieve this, upon MCTS backpropagation, we assign the same reward to any actions that lead to the same worklist sharding state. As a result, the MCTS expansion policy does not have to consider separate actions that lead to the same worklist state, effectively pruning the MCTS branching factor.


### 3.1.2 Meta-controller

Domain-specific enhancements reduce the search branching factor; however, the question of scaling to handle the multi-axis complexity remains. This work proposes using a meta-controller to focus the search effort on a single axis and a given objective. The meta-controller decomposes the global objective into a hierarchy of goals of the form (*axis, objective*). The order of goals influences the overall partitioning objective and reflects the expert’s preferences.


The meta-controller allocates a fixed budget to each goal and guides the search to focus on one of them. Once a goal exhausts its allocated budget, the meta-controller analyses the best-found partitioning strategy and replaces the model with the partitioned model if it finds it would improve on the goal’s objective. After that, the search focus on the next goal. This design enables the search to discover composite strategies by independently searching for partitioning strategies on each axis, only considering a single objective. Exposing these scripted goals to the users allows them to encode bespoke requirements, such as finer control over the trade-off between runtime and memory.

### 3.1.3 MCTS enhancements evaluation

This section evaluates how the enhancements described above enable Automap to work efficiently in a multi-axis setting. Transformers are well understood and studied in the DNN scaling literature [16, 10] and existing frameworks, such as DeepSpeed [24] Megatron-LM [28, 20], implement well-known partitioning strategies for their partition (e.g. Megatron [28, 20] and ZeRO [23]). A success metric for our work is to automatically discover Batch and Megatron model parallelism (BP and MP, respectively), together with ZeRO-3 partition, given a Transformer model.



(a) The odds that search discovers the composite sharding strategy "BP+MT+ZeRO3" in 25 seeds.



(b) The enhancements impact on the search throughput. We report the average for 25 different seeds. Higher throughput is better.

Figure 3: MCTS enhancement that enables multi-axis search. ST: the state representation enhancement, AG: action grouping, MC: meta-controller. + indicates the enhancement is enabled. Note that it is impossible to do action grouping without partitioned arguments as the state.

Figure 3a shows the benefit of using the MCTS enhancements in improving the odds of discovering the composite sharding strategy of BP+MT+ZeRO3. In all cases, only with the aid of the meta-controller was the search able to discover the composite strategy. Enabling each enhancement improved the odds of discovering the strategy, especially in constrained budget scenarios.

The search runs as part of the researcher development workflow and needs to be quick. In production, our search budget is, on average,  $\approx 1000 - 2000$  trajectories. As the search is constrained for trajectories, the enhancements must not negatively impact its throughput. Figure 3b evaluates their impact. Using the meta-controller improves the search throughput significantly as it allows the rewrite engine to skip checking the validity of partitioning actions on the inactive axis as the meta-controller ensures the search is acting only on a single axis. The small reduction in throughput because of the state and action grouping enhancement is mitigated by the fact that they aid in discovering complex solutions sooner and are more robust, as shown in Figure 3a.

Next, we investigate how the meta-controller utilizes its budget and allocate a large budget (4000 trajectories) to it. Figure 4 shows the trajectories needed to achieve each expert-level partitioning strategy. First, the meta-controller divides the maximum budget evenly between all the goals, which has the side-effect that the search spends too many cycles on one goal; the search discovered BP's strategy in 50 trajectories, and it continued exploring until it exhausted its budget. Similarly, the MT goal was achieved in the 250 trajectories. After that, however, it had to wait for the first goal to exhaust its budget before continuing. The shaded bar shows the opportunity to improve the search ability to finish quicker if it utilizes its budget smarter. We leave this as a discussion point on removing the need for budget allocation or providing a smarter budget estimate for each axis. Another opportunity for future work is the ability to discover these high-level goals automatically from data.




Figure 4: The trajectories needed to discover expert partitioning strategies automatically. Validated by analyzing the collectives reported in Table 1.

Table 1: The number of SPMD collectives of the discovered partitioning strategies that correspond to manual expert partitioning.

Strategy	AllGather	AllReduce	ReduceScatter
BP	0	387	0
BP+MT	0	483	0
BP+MT+ZeRO3	578	97	386

## 4 Evaluation

**Case-studies.** We evaluated our multi-axis search on three different model architectures:

- A 24-layer transformer model similar to GPT-3 [2] with input batch size of 16,  $d_{head} = 96$ ,  $n_{head} = 24$ ,  $d_{model} = 1024$ . Training this model has an estimated 280GB memory requirement making it infeasible to fit on any single device.
- A Graph Network Simulator (GNS) [26] model used in molecular structure prediction [7], configured with 24 layers of message passing, 1024 MLP hidden sizes as well as 512 latent sizes, with up to 10 molecule graphs per input batch (dynamically batched using graph [6]).
- A UNet of the reverse process in a diffusion model [9], with a 32 batch size. It uses 9 residual blocks for the down-sampling convolutions and up-sampling one, with two residual blocks in-between and an attention layer at  $16 \times 8 \times 4$  resolution with  $n_{heads} = 16$ .

**Hardware.** The transformer experiments were conducted on 32 TPU v3 cores laid out in an 8x4 2D mesh. The UNet and GNS models are smaller and conducted on 8 TPU v3 cores in a 2x4 2D mesh. The hardware runtime was averaged over ten measurements following one warmup step.

**Meta-controller settings.** A key design decision for the meta-controller is its ability to translate user requirements to meaningful search targets. At a high level, users want programs to fit into device memory and run as fast as possible. Here, we experiment with how different manually specified goals help achieve partitioning objectives. Eventually, we expect the meta-controller to handle the sequencing of objectives automatically. We tested four predefined goals:

- RT\_MEM\_ALL: Improve runtime on each axis then reduce the memory on each axis.
- RT1\_RT2\_MEM1: Improve runtime on each axis and then aim to reduce the memory on the batch axis. This goal is based on the same rationale behind a ZeRO3-like partitioning.
- RT1\_RT2\_MEM2: Similar to the previous but reduces memory usage on the model axis.
- RT\_MP\_ALL: Improve runtime on each axis while applying a progressively heavier memory penalty if the model does not fit in memory.
- No controller: No meta-controller, the search runs with a cost function that estimates runtime and a memory penalty term. The penalty term applies a linear cost based on exceeding the available memory to discourage partitions that do not fit.

As a simplification, we allowed the search a budget of 4000 trajectories with our compiler (but no actual hardware compilation), and the search results were compiled and tested on hardware.

In the remainder of this section, we evaluate the different meta-controller strategies and compare them to the known expert partitioning baselines. Figure 5b shows the runtime-memory plot once the model has been compiled and run on hardware. There is an obvious trade-off between the model’s runtime and peak memory, with the leftmost and bottom-most being the best points of that trade-off. These points drive the user requirement in designing the goals for the meta-controller. For example, a user who needs a ZeRO3-like partitioning strategy that provides memory reduction while maintaining a high training throughput RT1\_RT2\_MEM1 is an ideal strategy. However, a user constrained for memory would opt for a strategy focused on memory reduction. For example, in our initial experiments, we ran the same model with a 2x2 TPU topology (which has 16 devices). All strategies failed to fit on-device memory (even ZeRO3 could only get the model down to 17GB > 16GB available on TPU v3); all strategies except for the RT\_MEM\_ALL that was able to bring it down to 11GB by focusing on aggressive memory sharding, and still achieved an average of 1-second runtime (compared to the initial model estimate of 8 seconds).

To further validate the search performance, we evaluated it on other models of very different architecture: UNet and GNS. UNet does not have a known best expert partitioning strategy. Therefore, we used naive batch parallelism with ZeRO3 as a baseline. For GNS, we use edge-sharding [8] as a baseline. Edge-sharding shards the edges into sub-graphs on each axis, such that the computation is run concurrently. In particular, the original single-axis Automap has been used successfully to partition GNS-style models [36] via edge sharding.


We can see in Figure 5c that GNS strategies focused on aggressively reducing the memory or runtime have a significant advantage over edge-sharding. This illustrates the existence of non-trivial

partitioning strategies that could partition the graph into smaller sub-graphs to run concurrently, and each device only needs to store a smaller subset of the parameter and optimizer states. Similarly, in Figure 5d, UNet exhibits an interesting point where there is a better policy than all other strategies. This evaluation demonstrates that there are models, which are not being actively researched in the scaling literature, that have the potential to train twice as fast with half the resources, even when compared to strategies that are thought to be very effective (batch parallelism and ZeRO3).


**Action sequence length.** The depth of the decision tree is shortened by using the PartIR compiler stack and its powerful sharding propagation technique and grouping of arguments. To discover the BP+MT+ZeRO3 transformer sharding, the agent had to perform 13 actions; in contrast, to discover our unique UNet solution, the search had to perform  $30 \approx 40$  actions. Deciding the ideal depth of the tree leads to faster convergence as the search does not have to go deeper than it should and allows us to find unique solutions for models that require a higher-than-average sequence of sharding decisions. This decision can be learned or integrated through the powerful compiler analyzer toolchain.

(a) The reference estimated initial non-partitioned models performance.


Mode	Memory (GB)	Runtime (ms)
Transformer24	284.0	8740.0
GNS	12.592	521.48
UNet	8.761	84.9



(b) A 24-layer transformer model. BP runs out of memory and is not in the figure.



(c) A GNS model.



(d) A UNet like-model.

Figure 5: A runtime and memory performance for the models after partitioning using different partitioning strategies. The lower values are better. Shaded points are from different search seeds, with the median search result in a bolder color.

## 5 Discussion and outlook

This work explores the idea of incorporating several optimization objectives (goals) to discover composite partitioning strategies and proposes to improve MCTS efficiency through compiler analysis. There are several simplifications made in this work that are worth revisiting. Budget estimation for each goal is non-trivial, but further compiler analysis to approximate the cardinality of the independent partitioning actions can be used to bound both episode lengths and budget. Moreover, balancing local and global optimization goals is challenging and may require re-entrant search [21]. Finally, searching over heterogeneous SPMD meshes or temporal/pipeline strategies has not been explored yet but opens additional opportunities for meta-controller improvements. We also point out that leveraging compiler analysis in search is not free, and cost increases with model scale. In future work, we plan to analyze the trade-offs between searching and analyzing in more depth.

## References

- [1] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [2] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *CoRR*, abs/2005.14165, 2020. URL <https://arxiv.org/abs/2005.14165>.
- [3] Cameron Browne, Edward Powley, Daniel Whitehouse, Simon Lucas, Peter I. Cowling, Stephen Tavener, Diego Perez, Spyridon Samothrakis, Simon Colton, and et al. A survey of monte carlo tree search methods. *IEEE TRANSACTIONS ON COMPUTATIONAL INTELLIGENCE AND AI*, 2012.
- [4] Benjamin E Childs, James H Brodeur, and Levente Kocsis. Transpositions and move groups in monte carlo tree search. In *2008 IEEE Symposium On Computational Intelligence and Games*, pages 389–395. IEEE, 2008.
- [5] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways, 2022. URL <https://arxiv.org/abs/2204.02311>.
- [6] Jonathan Godwin\*, Thomas Keck\*, Peter Battaglia, Victor Bapst, Thomas Kipf, Yujia Li, Kimberly Stachenfeld, Petar Veličković, and Alvaro Sanchez-Gonzalez. Jraph: A library for graph neural networks in jax., 2020. URL <http://github.com/deepmind/jraph>.
- [7] Jonathan Godwin, Michael Schaarschmidt, Alexander L Gaunt, Alvaro Sanchez-Gonzalez, Yulia Rubanova, Petar Veličković, James Kirkpatrick, and Peter Battaglia. Simple GNN regularisation for 3d molecular property prediction and beyond. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=1wVvweK3oIb>.
- [8] Yilin He, Chaojie Wang, Hao Zhang, Bo Chen, and Mingyuan Zhou. Edge partition modulated graph convolutional networks, 2022. URL <https://openreview.net/forum?id=ET1UA0YeU42>.
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [10] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models, 2022. URL <https://arxiv.org/abs/2203.15556>.



- [11] Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Xu Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V. Le, Yonghui Wu, and Zhifeng Chen. Gpipe: Efficient training of giant neural networks using pipeline parallelism. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 103–112, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/093f65e080a295f8076b1c5722a46aa2-Abstract.html>.
- [12] Zhihao Jia, Oded Padon, James Thomas, Todd Warszawski, Matei Zaharia, and Alex Aiken. Taso: Optimizing deep learning computation with automatic generation of graph substitutions. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP ’19*, page 47–62, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450368735. doi: 10.1145/3341301.3359630. URL <https://doi.org/10.1145/3341301.3359630>.
- [13] Zhihao Jia, Matei Zaharia, and Alex Aiken. Beyond data and model parallelism for deep neural networks. In Ameet Talwalkar, Virginia Smith, and Matei Zaharia, editors, *Proceedings of Machine Learning and Systems 2019, MLSys 2019, Stanford, CA, USA, March 31 - April 2, 2019*. mlsys.org, 2019. URL <https://proceedings.mlsys.org/book/265.pdf>.
- [14] Norman P. Jouppi, Cliff Young, Nishant Patil, David A. Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, Richard C. Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. In-datacenter performance analysis of a tensor processing unit. *CoRR*, abs/1704.04760, 2017. URL <http://arxiv.org/abs/1704.04760>.
- [15] Norman P. Jouppi, Doe Hyun Yoon, Matthew Ashcraft, Mark Gottscho, Thomas B. Jablin, George Kurian, James Laudon, Sheng Li, Peter Ma, Xiaoyu Ma, Thomas Norrie, Nishant Patil, Sushma Prasad, Cliff Young, Zongwei Zhou, and David Patterson. Ten lessons from three generations shaped google’s tpuv4i : Industrial product. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pages 1–14, 2021. doi: 10.1109/ISCA52012.2021.00010.
- [16] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *CoRR*, abs/2001.08361, 2020. URL <https://arxiv.org/abs/2001.08361>.
- [17] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. Mlir: Scaling compiler infrastructure for domain specific computation. In *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 2–14, 2021. doi: 10.1109/CGO51591.2021.9370308.
- [18] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *CoRR*, abs/2006.16668, 2020. URL <https://arxiv.org/abs/2006.16668>.
- [19] Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R. Devanur, Gregory R. Ganger, Phillip B. Gibbons, and Matei Zaharia. Pipedream: generalized pipeline parallelism for DNN training. In Tim Brecht and Carey Williamson, editors, *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019, Huntsville, ON,*

- Canada, October 27-30, 2019, pages 1–15. ACM, 2019. doi: 10.1145/3341301.3359646. URL <https://doi.org/10.1145/3341301.3359646>.
- [20] Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. Efficient large-scale language model training on GPU clusters. *CoRR*, abs/2104.04473, 2021. URL <https://arxiv.org/abs/2104.04473>.
- [21] Phitchaya Mangpo Phothilimthana, Amit Sabne, Nikhil Sarda, Karthik Srinivasa Murthy, Yanqi Zhou, Christof Angermueller, Mike Burrows, Sudip Roy, Ketan Mandke, Reza Farahani, Yu Emma Wang, Berkin Ilbeyi, Blake Hechtman, Bjarke Roune, Shen Wang, Yuanzhong Xu, and Samuel J. Kaufman. A flexible approach to autotuning multi-pass machine learning compilers. In *2021 30th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 1–16, 2021. doi: 10.1109/PACT52795.2021.00008.
- [22] Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, H. Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Maribeth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu, Erich Elsen, Siddhant M. Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsimpoukelli, Nikolai Grigorev, Doug Fritz, Thibault Sottiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d’Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake A. Hechtman, Laura Weidinger, Iason Gabriel, William S. Isaac, Edward Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu, and Geoffrey Irving. Scaling language models: Methods, analysis & insights from training gopher. *CoRR*, abs/2112.11446, 2021. URL <https://arxiv.org/abs/2112.11446>.
- [23] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimization towards training A trillion parameter models. *CoRR*, abs/1910.02054, 2019. URL <http://arxiv.org/abs/1910.02054>.
- [24] Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD ’20*, page 3505–3506, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. doi: 10.1145/3394486.3406703. URL <https://doi.org/10.1145/3394486.3406703>.
- [25] Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. Zero-offload: Democratizing billion-scale model training. *CoRR*, abs/2101.06840, 2021. URL <https://arxiv.org/abs/2101.06840>.
- [26] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning*, pages 8459–8468. PMLR, 2020.
- [27] Michael Schaarschmidt, Dominik Grewe, Dimitrios Vytiniotis, Adam Paszke, Georg Stefan Schmid, Tamara Norman, James Molloy, Jonathan Godwin, Norman Alexander Rink, Vinod Nair, et al. Automap: Towards ergonomic automated parallelism for ml models. *arXiv preprint arXiv:2112.02958*, 2021.
- [28] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *CoRR*, abs/1909.08053, 2019. URL <http://arxiv.org/abs/1909.08053>.

- [29] Colin Unger, Zhihao Jia, Wei Wu, Sina Lin, Mandeep Baines, Carlos Efrain Quintero Narvaez, Vinay Ramakrishnaiah, Nirmal Prajapati, Pat McCormick, Jamaludin Mohd-Yusof, Xi Luo, Dheevatsa Mudigere, Jongsoo Park, Misha Smelyanskiy, and Alex Aiken. Unity: Accelerating DNN training through joint optimization of algebraic transformations and parallelization. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 267–284, Carlsbad, CA, July 2022. USENIX Association. ISBN 978-1-939133-28-1. URL <https://www.usenix.org/conference/osdi22/presentation/unger>.
- [30] Dimitrios Vytiniotis. Declarative abstractions for tensor program partitioning. In *Proceedings of the 22nd International Symposium on Principles and Practice of Declarative Programming*, PPDP '20, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450388214. doi: 10.1145/3414080.3414105. URL <https://doi.org/10.1145/3414080.3414105>.
- [31] Siyu Wang, Yi Rong, Shiqing Fan, Zhen Zheng, Lansong Diao, Guoping Long, Jun Yang, Xiaoyong Liu, and Wei Lin. Auto-map: A DQN framework for exploring distributed execution plans for DNN workloads. *CoRR*, abs/2007.04069, 2020. URL <https://arxiv.org/abs/2007.04069>.
- [32] Xinfeng Xie, Prakash Prabhu, Ulysse Beaugnon, Phitchaya Mangpo Phothilimthana, Sudip Roy, Azalia Mirhoseini, Eugene Brevdo, James Laudon, and Yanqi Zhou. A transferable approach for partitioning machine learning models on multi-chip-modules. In Diana Marculescu, Yuejie Chi, and Carole-Jean Wu, editors, *Proceedings of Machine Learning and Systems 2022, MLSys 2022, Santa Clara, CA, USA, August 29 - September 1, 2022*. mlsys.org, 2022. URL <https://proceedings.mlsys.org/paper/2022/hash/7f1de29e6da19d22b51c68001e7e0e54-Abstract.html>.
- [33] Yuanzhong Xu, HyoukJoong Lee, Dehao Chen, Hongjun Choi, Blake Hechtman, and Shibo Wang. Automatic cross-replica sharding of weight update in data-parallel training, 2020. URL <https://arxiv.org/abs/2004.13336>.
- [34] Yuanzhong Xu, HyoukJoong Lee, Dehao Chen, Blake A. Hechtman, Yanping Huang, Rahul Joshi, Maxim Krikun, Dmitry Lepikhin, Andy Ly, Marcello Maggioni, Ruoming Pang, Noam Shazeer, Shibo Wang, Tao Wang, Yonghui Wu, and Zhifeng Chen. GSPMD: general and scalable parallelization for ML computation graphs. *CoRR*, abs/2105.04663, 2021. URL <https://arxiv.org/abs/2105.04663>.
- [35] Bowen Yang, Jian Zhang, Jonathan Li, Christopher Ré, Christopher R. Aberger, and Christopher De Sa. Pipemare: Asynchronous pipeline parallel DNN training. *CoRR*, abs/1910.05124, 2019. URL <http://arxiv.org/abs/1910.05124>.
- [36] Sheheryar Zaidi, Michael Schaarschmidt, James Martens, Hyunjik Kim, Yee Whye Teh, Alvaro Sanchez-Gonzalez, Peter Battaglia, Razvan Pascanu, and Jonathan Godwin. Pre-training via denoising for molecular property prediction, 2022. URL <https://arxiv.org/abs/2206.00133>.
- [37] Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yuanzhong Xu, Danyang Zhuo, Joseph E. Gonzalez, and Ion Stoica. Alpa: Automating inter- and intra-operator parallelism for distributed deep learning. *CoRR*, abs/2201.12023, 2022. URL <https://arxiv.org/abs/2201.12023>.