Hyper: Distributed Cloud Processing for Large-Scale Deep Learning Tasks

Davit Buniatyan

Snark AI, Menlo Park, CA 94025

Princeton University, Princeton, NJ 08544

davit@snark.ai

Abstract—**Training and deploying deep learning models in real-world applications require processing large amounts of data. This is a challenging task when the amount of data grows to a hundred terabytes, or even, petabyte-scale. We introduce a hybrid distributed cloud framework with a unified view to multiple clouds and an on-premise infrastructure for processing tasks using both CPU and GPU compute instances at scale. The system implements a distributed file system and failure-tolerant task processing scheduler, independent of the language and Deep Learning framework used. It allows to utilize unstable cheap resources on the cloud to significantly reduce costs. We demonstrate the scalability of the framework on running pre-processing, distributed training, hyperparameter search and large-scale inference tasks utilizing 10,000 CPU cores and 300 GPU instances with overall processing power of 30 petaflops.**[1][2]**

Index Terms—**Deep Learning, Cloud Computing, Distributed Systems**

I. INTRODUCTION

Deep Learning (DL) based models have outperformed manual feature engineered algorithms in a wide range of domains including computer vision, natural language processing, audio processing [8], [15], [22], [23], [25], [29], [41]. The amount of labelled data required for training production-ready models achieves a terabyte-scale [14], [37]. The unlabelled data to execute those models reaches a petabyte-scale [5], [7], [42]. Such computational resources are on-demand available on clouds such as Amazon Web Services (AWS), Google Cloud Platform (GCP) or Azure.

Modern Deep Learning frameworks such as PyTorch [28], Tensorflow [6], MXNet [12], Theano [9] and Caffe [19] are well-positioned for training and deploying models on single multicore machines with multiple GPUs or TPUs. Training state-of-the-art models on terascale data often takes weeks or months to converge. To make training faster, DL frameworks such as PyTorch or Tensorflow recently introduced synchronous and asynchronous distributed training methods to achieve an almost linear speedup with respect to the number of nodes.

As the number of nodes in a distributed cluster grows, problems such as provisioning, orchestrating, fault-tolerance, distribution data management, task planning, and execution arises. For executing Big Data workloads, several widely accepted synchronous parallel processing systems have been introduced such as MapReduce [13], Apache Spark [40] and Dryad [18]. Additionally, task-parallel frameworks are getting increased usage such as Dask [31] or CIEL [27]. They provide fine-grained task management. These frameworks are very efficient for ETL tasks but lack native deep learning support. A recent introduction to the family is Ray, which implements a dynamic task-parallel framework that is suited for deep learning and reinforcement learning [26].

To manage High-Performance Computing (HPC) workloads, container technology has recently become well suite choice for packaging environment libraries [39]. Frameworks such as Kubernetes enable massive distribution of stateless applications packaged in a container on a cluster of nodes in fault-tolerant way [10]. However, it still lacks an efficiently distributed data management unit, workflow scheduling system, and support for stateful applications such as model training. Packages such as KubeFlow and Argo attempt to extend Kubernetes to support the implementation of machine learning pipelines [17].

For data-intensive workloads, a variety of distributed file storage systems have been introduced [32] such as NFS [34] or HDFS [35]. NFS-based file systems significantly decrease multi-read speed and lower bound the computing speed. They often do not scale on multi-write scenarios. There is always a trade-off between latency and scalability. For web-based applications, cloud providers offer an object storage solution that has high scalability but suffers from low-latency memory-intensive operations.

We introduce a hybrid distributed cloud framework with a unified view of multiple clouds and on-premise infrastructure for processing tasks. We have made the following contributions

- We design and build a distributed framework that unifies preprocessing, training, hyperparameter search and deploying large scale applications on the cloud.
- To achieve scalability we design and deploy a distributed file system that has near-zero delay for deep learning jobs in comparison to downloading the data locally on the machine with similar performance.
- Provide fault-tolerance with respect to computational node failures and support utilization of unstable cheap resources on the cloud.

[1] Demo available at https://lab.snark.ai

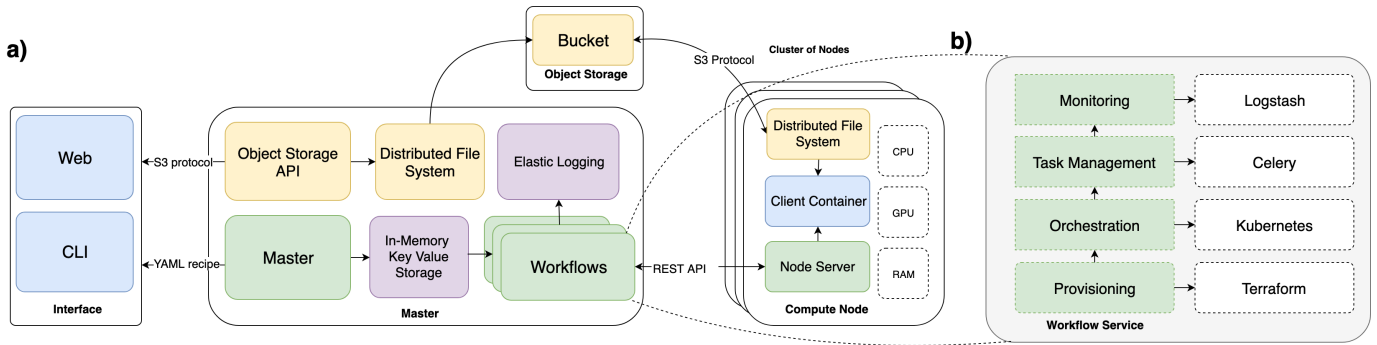[2] Documentation available at https://docs.snark.ai

Fig. 1. System Architecture a) Interface uploads the training data, source code and YAML recipe to the Master Node. Source files are chunked and uploaded to Object Storage. The Recipe is parsed to create a computational graph in in-memory Key-Value Storage. For each workflow, the cluster is created in the cloud. Each computational node has a Node Server that handles the management of the node and executing client containers. b) Workflow has four main stages. Provisioning the infrastructure, orchestrating nodes, executing tasks and monitoring workers.

- Abstract away cloud infrastructure setup and provisioning from the user with native cloud integration.

II. COMPUTATIONAL MODEL

In this section, we discuss design decisions made for the system and user interface to specify the computation workload.

A. Computational Workflows

Workflow is a directed acyclic graph consisting of *Experiment* nodes and their dependency as edges. Single $Experiment$ contains multiple $Tasks$. $Tasks$ within the same experiment execute the same command with different arguments. Arguments can be templated for efficient parameter space definitions. Each *Experiment* has an associated container that is deployed on all computational workers.

Task is the execution unit, which encapsulates a process. Each *Task* has assigned *Node*, which represents the computation worker. Single *Node* might execute multiple *Tasks*. The number of nodes available corresponds to the number of workers inside the cluster.

B. Interface

$Workflows$ are specified using code-as-infrastructure interface defined in YAML recipes as seen in Fig. **??**. The recipe is parsed by the server and translated into a directed acyclic graph of experiments. The interface lets users specify the environment, hardware settings, number of workers, parameters and parameterized commands. The user can interface the system through CLI or Web UI.

C. Parameters

The user can specify the list of parameters that can be inserted into command arguments during execution. Parameters can be sampled from a discrete class or continuous range. To compute parameters for each *Task*, the algorithm generates the Cartesian product of all discrete parameters and samples from the set n times with minimal repetition. n is defined to be the number of samples from a recipe. Then, it samples n times from each continuous parameter range and randomly matches

with discrete sampled parameters. This is necessary to support both hyper-parameter search and inference with grid iterators.

III. SYSTEM OVERVIEW

The system receives data, chunks it and stores it in object storage. The recipe is submitted to deploy the deep learning workflow on a cluster of nodes, which mounts the distributed file system.

A. Distributed Data Management

In order to be framework agnostic and require no further modification of the client program, we chunk the file system itself and store it in object storage provided by the cloud vendor (e.g. AWS S3). The system implementation is similar to closed source ObjectiveFS [1] and tuned for deep learning tasks. The distributed file-system wraps POSIX API and acts as a middle layer with chunking, caching and state synchronization mechanisms across all nodes. When the program queries the file system for a specific file, the integration layer checks which chunk contains the file to download. In the next query, the file system can check if the existing chunk contains the next required file before fetching it from the cloud. Within the program's context, files that are stored in remote chunked object storage appear to be local files. Any deep learning application without further modification will take advantage of the highly scalable storage.

Deep learning frameworks such as PyTorch and Tensor-Flow natively support asynchronous data fetching from the local storage to the GPU using data loaders. Often the deep learning training iteration is bounded by the compute cycles on GPUs. If one combines the distributed remote storage and asynchronous data fetching, the training speed is almost the same as if the data was stored locally on the machine with respect to the following constraints.

- Network locality - We assume that computational nodes that access the data are physically located near object storage servers.
- Chunksize - well chosen with respect to latency to maximize the throughput as shown in Fig. 2. It should be in the range 12-100MB.
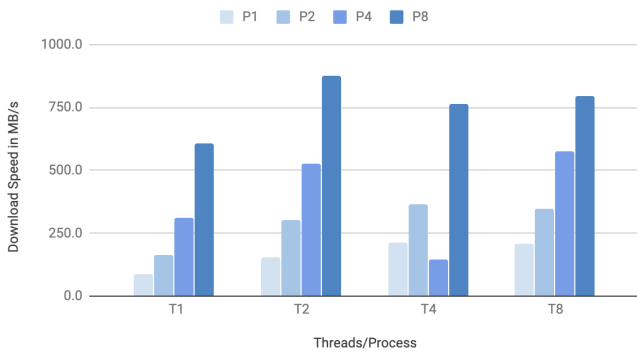
Fig. 2. Hyper File System on a single machine (AWS p3.2xlarge) can achieve up to 875 MB/s download speed with multithreading T and multiprocessing P enabled. The batch size was chosen smaller for large models to fit in the GPU RAM.

The suggested file system can leverage the scalability of the object storage and provide data access to the cluster of nodes with almost native speed for deep learning jobs as shown in Fig. 3.

B. Cloud Infrastructure

Distributed frameworks such as Spark, provide a functional ecosystem for computational tasks. End-users should only specify the program and apply it to large datasets. User is limited to only using supported libraries abstracted in MapReduce framework. State-of-the-art libraries are out of reach.

Provisioning: When constraints of the system does allow arbitrary package support, then the whole environment and necessary packages should be transported to the computational node. We use container technology to bake the necessary libraries. Compute nodes need to have docker support to execute arbitrary containers and Nvidia CUDA [21] libraries for processing deep learning operations on the GPU.

Orchestration: Due to its generic feature of container technology, the Virtual Machine (VM) images necessary to run containers can be based on any Unix based operating systems including CoreOS, Ubuntu and CentOS. VM images are built only once and stored in the cloud. It acts as a proxy to execute custom specified containers. The user can specify the container from the public repository. After cloud instances are provisioned, each instance downloads a client container. This mechanism allows supporting any framework, library or package without constraints. We also cache frequently used containers such as Tensorflow, Pytorch, Jupyter directly inside VM images to reduce loading time. In addition to custom docker management, the system can offload container orchestration to managed Kubernetes [10].

Networking: For cloud infrastructure orchestration, we use Terraform, which provides code-as-infrastructure language for defining cloud resources. For each job execution, the system specifies a Virtual Private Network with Internet Gateway. It makes cluster nodes accessible inside the network for use cases that require state synchronization across nodes such as during distributed training. Alternatively one can use object storage as a parameter server to store the model without networking setup.

C. Implementation

As shown in Fig. 1, the architecture of the distributed framework consists of main components: Interface, Master and Node. Master is responsible for receiving the recipe of the pipeline, parsing and creating workflow objects including experiments and tasks. The objects are stored in-memory key-value cache Redis. As a backup alternative, the system stores the state into DynamoDB. Then, the master starts a new workflow service as an adjacent container to orchestrate and schedule tasks. During the orchestration process, each compute worker runs a node server that listens to commands executed by the workflow manager. Each node starts to pull the client specified container and mount the distributed file system. Once the node is ready, the workflow manager can execute the client's specified commands.

There are three types of logs that are collected into Elastic Logstash: client application logs, CPU/GPU utilization logs and operating system logs [16] .

In addition to the scheduling system, we deploy our own object storage layer for providing S3-like API [36]to client interface. When the data is uploaded to Minio server [4], files get chunked and stored on the distributed file system. Celery is used for asynchronous task management. The task management system is similar to Apache Airflow [3] and other workload management systems such as Splunk. It is different from frameworks such as Dask [31] , CIEL [27] or Ray [26] in terms of execution granularity.

D. Fault tolerance

In order to optimize cloud resource allocation, cloud vendors provide spots or preemptible instances. Those instances are usually 2 or 3 times cheaper but can be terminated anytime depending on the demand and the price per hour bid. For stateful long-lasting jobs, cost optimization is an attractive option, however, it requires additional compute logic implementation to recover the state.

Since the system already provides a distributed file system backed by remote object storage and a scheduling system, it becomes straightforward to implement a fault-tolerant system that can handle instability. When a node fails, the task with exact command arguments gets rescheduled on a different node. In training use cases, modern deep learning frameworks provide an easy interface to store and retrieve model states. Hence, the training can be continued without any additional code modifications.

IV. EVALUATION

To evaluate the system we run pre-processing, distributed training, hyper-parameter search and large scale inference using AWS CPU M5 family and GPU P3 family compute instances. We use AWS S3 to store file system chunks.

A. Preprocessing

To test the scalability of the system for ETL tasks, we set up a preprocessing experiment. 100 million text files from commoncrawl [5] dataset are uploaded to the distributed storage. The amount of data achieves 10TB. We specify the infrastructure to spin up 110 instances each with 96 CPU cores. The processing script takes 100,000 text files and transforms them into tfrecord files. During the transformation, spaCy [2] package is used for filtering, tokenizing and splitting paragraphs. We also enable spot instances to reduce costs and test fault tolerance.

B. Distributed Training

We defined a recipe for training object recognition model YoloV3 [30] using PyTorch [28] using Horovod [33] . Then, we uploaded COCO dataset [24] to the storage. The training script reads all images and labels from the defined path. Furthermore, parameters such as how many epochs, learning rate and what input image size are parameterized in the recipe. We also parametrize model-specific parameters such as a total number of classes, confidence threshold, nms threshold and iou threshold.

Nvidia K80 GPUs are slow to train the model. With a single line configuration change, we deployed the training on Nvidia V100 GPUs with spot enabled instances to reduce costs. The batch size for the training was accordingly modified. The cost would be $8.48/h instead of $0.95/h, but the training is 50x faster with 6x efficiency gain. We modified hyper-parameters and started another experiment with zero effort.

We also provide benchmarking of data streaming against storing data in local files while training data-intensive computer vision models as shown in Fig. 3 and Fig. 3.

C. Hyperparameter search

Gradient boosting machine, XGBoost [11] or LightGBM [20], is one of the best off-the-shelf machine learning solvers for tabular data, however training those models can be computationally very heavy and have a lot of parameters to tune. There are 12 parameters to tune for the tree booster. If you try 2 choices for each one, there will be 4096 different combinations. If each training takes 10 mins to complete, trying out all those 4096 combinations sequentially would take 28.4 days. Using our system, we made the experiments run in 10 minutes by linearly increasing the cluster size without source code modification.

D. Large-Scale Inference

For production-ready processing, we upload an ImageNet dataset [14] and split it into 300 folders. Each folder contains 1500 images. Using Hyper interface, we easily parallelized the inference execution of Yolo model to 300 GPU instances with the overall processing of 2 petaflops.
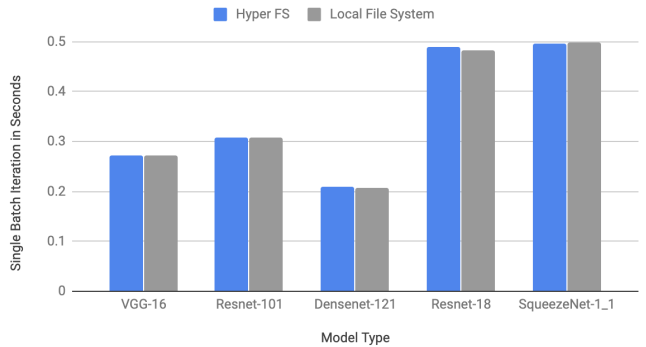

Fig. 3. Streaming data through Hyper File System while training a deep learning model is equivalent to reading data from the local file system
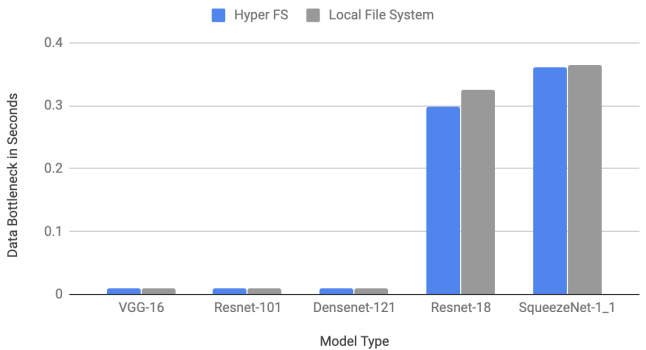

Fig. 4. In the asynchronous data loading figure, the first three models (VGG, Resnet101, and DenseNet) have no data bottleneck. The benchmarks have been produced on AWS using a p3.2xlarge machine with V100 GPU. The data is stored on S3 within the same region.

V. DISCUSSION

Building Hyper was inspired by processing a petascale amount of brain images for connectomics [42]. At the time, Map-Reduce based frameworks such as Spark and directed graph processing frameworks such as Dask or Ray have not been efficient in executing deep learning computes. Initially, we used AirFlow [3] , then used framework ChunkFlow [38]. Hyper is another generation of cloud-agnostic, data-agnostic distributed deep learning computed with scale. Since Hyper was released hundreds of people have used the platform and several companies have used it in production. Here we discuss some of the early feedback and limitations of the software.

Feedback When Hyper is used in production, it completely encapsulates infrastructure and is flexible enough to customize projects under user needs. The user only needs to upload the data and the source code to the distributed file storage through CLI or Web UI. Then submit the infrastructure as minimal code to describe the compute workflow. It provides interface to log results of hyperparameter search.

Ecosystem Hyper supports any open source machine learning libraries and frameworks. We can offer any pre-built algorithms and models that suit different business cases.

On-prem Deployment They system can deploy the software to any on-premise infrastructure with Kubernetes sup-

port.

Built-in Modeling Algorithms We support cutting-edge pre-built algorithms and models for a wide range of industry use-cases such as image recognition, speech to text, text to speech, anomaly detection, recommendation and personalization, and others.

Unlimited Persistent Data Lake Scalable and cost-efficient persistent central Data Lake that merges data from all of your different data sources. A single source of truth for a complete view of the datasets. 10x cheaper on the storage system compared to AWS EFS expenses with almost all native speed of data read/write.

Spot/Pre-emptible Instance Cost Savings AWS and GCP offer dynamic bidding for instances which are usually 3x cheaper than normal instances but also can be stopped at any time. Spot instance management layer makes it very easy to use spot instances and also enjoy the cost-saving.

Full lifecycle ML pipeline Drag and drop machine learning modules that include data ingestion modules, data ETL modules, built-in model training algorithm modules, hyper-parameter search modules, and model deployment modules. Combine the modules into a pipeline for a full lifecycle of your machine learning task.

Limitations Current limitations include low interactivity of the workflow. Once the workflow has been defined. Users are not able to modify the execution unless terminated or start a new one. Compared to Ray and Dask, Hyper modifies the surroundings of the application-defined by deep learning researcher - and has minimal impact on the code itself. Even though it brings several advantages for end-users, it also provides strict limitations on computing optimization and efficiency. Optimizations such as GPU full capacity utilization and inter-network communications are still handled by the user.

VI. Conclusion

Hyper provides a unified view of multiple clouds and on-premise infrastructure without requiring a team of DevOps engineers to save AI/ML-Ops time. It provides cloud cost savings and transparent compute resource utilization tracking. Hyper Storage solutions are significantly better for data-intensive deep learning tasks compared to cloud-native NFS-like offers and much more cost-efficient since they are backed by object storage.

Hyper enables data scientists to be highly efficient in machine learning and deep learning. It provides a framework for running experiments, collecting logs and comparing models including one-click Jupyter notebooks or Tensorboard graphs using Web UI or CLI. Data scientists can plug-and-play state-of-the-art deep learning models in Computer Vision, NLP, and other domains to kickstart their project. They can execute large-scale distributed training or batch processing jobs through a very simple interface without knowing about the infrastructure and define continuous workflows for automatic model training, validation, benchmarking and deployment.

Future work includes the development of various Bayesian optimization algorithms for hyper-parameter tuning of models, seamless integration with Kubernetes and interactive workflows.

VII. Acknowledgement

References

[1] ObjectiveFS: Shared S3 File System for EC2 instances and Servers . https://objectivefs.com/rg, 2013. [Online; accessed 31-May-2019].

[2] spaCy Industrial-Strength Natural Language Processing. https://github.com/explosion/spaCy, 2015. [Online; accessed 31-May-2019].

[3] Apache Airflow. http://airflow.incubator.apache.org, 2017. [Online; accessed 31-May-2019].

[4] MinIO high performance object storage server compatible with Amazon S3 API. https://github.com/minio/minio, 2018. [Online; accessed 31-May-2019].

[5] Common Crawl Dataset. https://commoncrawl.org, 2019. [Online; accessed 31-May-2019].

[6] 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.

[7] Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First m87 event horizon telescope results. iii. data processing and calibration. *The Astrophysical Journal Letters*, 875(1):L3, 2019.

[8] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

[9] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. Theano: a cpu and gpu math expression compiler. In *Proceedings of the Python for scientific computing conference (SciPy)*, volume 4. Austin, TX, 2010.

[10] David Bernstein. Containers and cloud: From lxc to docker to kubernetes. *IEEE Cloud Computing*, 1(3):81–84, 2014.

[11] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794. ACM, 2016.

[12] 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. *arXiv preprint arXiv:1512.01274*, 2015.

[13] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[14] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.

[15] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[16] Clinton Gormley and Zachary Tong. *Elasticsearch: The definitive guide: A distributed real-time search and analytics engine*. " O'Reilly Media, Inc.", 2015.

[17] Xinyuan Huang, Amit Kumar Saha, Debojyoti Dutta, and Ce Gao. Kubebench: A benchmarking platform for ml workloads. In *2018 First International Conference on Artificial Intelligence for Industries (AI4I)*, pages 73–76. IEEE, 2018.

[18] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. In *ACM SIGOPS operating systems review*, volume 41, pages 59–72. ACM, 2007.

[19] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 675–678. ACM, 2014.

[20] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3146–3154. Curran Associates, Inc., 2017.

[21] David Kirk et al. Nvidia cuda software and gpu parallel computing architecture. In *ISMM*, volume 7, pages 103–104, 2007.

[22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[23] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.

[24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

[25] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

[26] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. Ray: A distributed framework for emerging {AI} applications. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 561–577, 2018.

[27] Derek G Murray, Malte Schwarzkopf, Christopher Smowton, Steven Smith, Anil Madhavapeddy, and Steven Hand. Ciel: a universal execution engine for distributed data-flow computing. In *Proc. 8th ACM/USENIX Symposium on Networked Systems Design and Implementation*, pages 113–126, 2011.

[28] 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.

[29] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.

[30] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.

[31] Matthew Rocklin. Dask: Parallel computation with blocked algorithms and task scheduling. In *Proceedings of the 14th Python in Science Conference*, number 130-136. Citeseer, 2015.

[32] Mahadev Satyanarayanan. A survey of distributed file systems. *Annual Review of Computer Science*, 4(1):73–104, 1990.

[33] Alexander Sergeev and Mike Del Balso. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799*, 2018.

[34] Spencer Shepler, Brent Callaghan, David Robinson, Robert Thurlow, Carl Beame, Mike Eisler, and Dave Noveck. Network file system (nfs) version 4 protocol. Technical report, 2003.

[35] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, Robert Chansler, et al. The hadoop distributed file system. In *MSST*, volume 10, pages 1–10, 2010.

[36] Swaminathan Sivasubramanian. Amazon dynamodb: a seamlessly scalable non-relational database service. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 729–730. ACM, 2012.

[37] Baoyuan Wu, Weidong Chen, Yanbo Fan, Yong Zhang, Jinlong Hou, Junzhou Huang, Wei Liu, and Tong Zhang. Tencent ml-images: A large-scale multi-label image database for visual representation learning. *arXiv preprint arXiv:1901.01703*, 2019.

[38] Jingpeng Wu, William M Silversmith, and H Sebastian Seung. Chunkflow: Distributed hybrid cloud processing of large 3d images by convolutional nets. *arXiv preprint arXiv:1904.10489*, 2019.

[39] Miguel G Xavier, Marcelo V Neves, Fabio D Rossi, Tiago C Ferreto, Timoteo Lange, and Cesar AF De Rose. Performance evaluation of container-based virtualization for high performance computing environments. In *2013 21st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*, pages 233–240. IEEE, 2013.

[40] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. *HotCloud*, 10(10-10):95, 2010.

[41] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.

[42] Zhihao Zheng, J Scott Lauritzen, Eric Perlman, Camenzind G Robinson, Matthew Nichols, Daniel Milkie, Omar Torrens, John Price, Corey B Fisher, Nadiya Sharifi, et al. A complete electron microscopy volume of the brain of adult drosophila melanogaster. *Cell*, 174(3):730–743, 2018.