

---

# MLPERF INFERENCE BENCHMARK

---

Vijay Janapa Reddi<sup>1</sup> Christine Cheng<sup>2</sup> David Kanter<sup>3</sup> Peter Mattson<sup>4</sup> Guenther Schmuelling<sup>5</sup>  
Carole-Jean Wu<sup>6</sup> Brian Anderson<sup>4</sup> Maximilien Breughe<sup>7</sup> Mark Charlebois<sup>8</sup> William Chou<sup>8</sup>  
Ramesh Chukka<sup>2</sup> Cody Coleman<sup>9</sup> Sam Davis<sup>10</sup> Pan Deng<sup>11</sup> Greg Damos<sup>12</sup> Jared Duke<sup>4</sup> Dave Fick<sup>13</sup>  
J. Scott Gardner<sup>14</sup> Itay Hubara<sup>15</sup> Sachin Idgunji<sup>7</sup> Thomas B. Jablin<sup>4</sup> Jeff Jiao<sup>16</sup> Tom St. John<sup>17</sup>  
Pankaj Kanwar<sup>4</sup> David Lee<sup>18</sup> Jeffery Liao<sup>19</sup> Anton Lokhmotov<sup>20</sup> Francisco Massa<sup>6</sup> Peng Meng<sup>11</sup>  
Paulius Micikevicius<sup>7</sup> Colin Osborne<sup>21</sup> Gennady Pekhimenko<sup>22</sup> Arun Tejusve Raghunath Rajan<sup>2</sup>  
Dilip Sequeira<sup>7</sup> Ashish Sirasao<sup>23</sup> Fei Sun<sup>24</sup> Hanlin Tang<sup>2</sup> Michael Thomson<sup>25</sup> Frank Wei<sup>26</sup> Ephrem Wu<sup>23</sup>  
Lingjie Xu<sup>26</sup> Koichi Yamada<sup>2</sup> Bing Yu<sup>18</sup> George Yuan<sup>7</sup> Aaron Zhong<sup>16</sup> Peizhao Zhang<sup>6</sup> Yuchen Zhou<sup>27</sup>

Machine-learning (ML) hardware and software system demand is burgeoning. Driven by ML applications, the number of different ML inference systems has exploded. Over 100 organizations are building ML inference chips, and the systems that incorporate existing models span at least three orders of magnitude in power consumption and four orders of magnitude in performance; they range from embedded devices to data-center solutions. Fueling the hardware are a dozen or more software frameworks and libraries. The myriad combinations of ML hardware and ML software make assessing ML-system performance in an architecture-neutral, representative, and reproducible manner challenging. There is a clear need for industry-wide standard ML benchmarking and evaluation criteria. MLPerf Inference answers that call. Driven by more than 30 organizations as well as more than 200 ML engineers and practitioners, MLPerf implements a set of rules and practices to ensure comparability across systems with wildly differing architectures. In this paper, we present the method and design principles of the initial MLPerf Inference release. The first call for submissions garnered more than 600 inference-performance measurements from 14 organizations, representing over 30 systems that show a range of capabilities.

## 1 INTRODUCTION

Machine learning (ML) powers a variety of applications from computer vision (He et al., 2016; Goodfellow et al., 2014; Liu et al., 2016; Krizhevsky et al., 2012) and natural-language processing (Vaswani et al., 2017; Devlin et al., 2018) to self-driving cars (Xu et al., 2018; Badrinarayanan et al., 2017) and autonomous robotics (Levine et al., 2018). These applications are deployed at large scale and require substantial investment to optimize inference performance. Although training of ML models has been a development bottleneck and a considerable expense (Amodei & Hernandez, 2018), inference has become a critical workload, since models can serve as many as 200 trillion queries and perform over 6 billion translations a day (Lee et al., 2019b).

To address these growing computational demands, hardware, software, and system developers have focused on inference performance for a variety of use cases by designing optimized ML hardware and software systems. Estimates indicate that over 100 companies are producing or are on the verge of producing optimized inference chips. By comparison, only about 20 companies target training.

Each system takes a unique approach to inference and presents a trade-off between latency, throughput, power, and model quality. For example, quantization and reduced precision are powerful techniques for improving inference latency, throughput, and power efficiency at the expense of accuracy (Han et al., 2015; 2016). After training with floating-point numbers, compressing model weights enables better performance by decreasing memory-bandwidth requirements and increasing computational throughput (e.g., by using wider vectors). Similarly, many weights can be removed to boost sparsity, which can reduce the memory footprint and the number of operations (Han et al., 2015; Molchanov et al., 2016; Li et al., 2016). Support for these techniques varies among systems, however, and these optimizations can drastically reduce final model quality. Hence, the field needs an ML inference benchmark that can quantify these trade-offs in an architecturally neutral, representative, and reproducible manner.

---

<sup>1</sup>Harvard University <sup>2</sup>Intel <sup>3</sup>Real World Insights <sup>4</sup>Google  
<sup>5</sup>Microsoft <sup>6</sup>Facebook <sup>7</sup>NVIDIA <sup>8</sup>Qualcomm <sup>9</sup>Stanford University  
<sup>10</sup>Myrtle <sup>11</sup>Tencent <sup>12</sup>Landing AI <sup>13</sup>Mythic <sup>14</sup>Advantage Engineering  
<sup>15</sup>Habana Labs <sup>16</sup>Alibaba T-Head <sup>17</sup>Tesla <sup>18</sup>MediaTek  
<sup>19</sup>Synopsys <sup>20</sup>dividiti <sup>21</sup>Arm <sup>22</sup>University of Toronto <sup>23</sup>Xilinx  
<sup>24</sup>Alibaba (formerly Facebook) <sup>25</sup>Centaur Technology <sup>26</sup>Alibaba Cloud  
<sup>27</sup>General Motors. MLPerf Inference is the product of individuals from these organizations who led the benchmarking effort and of submitters who produced the first set of benchmark results. Both groups are necessary to create a successful industry benchmark. We credit the submitters and their organizations in the acknowledgments. Send correspondence to [vj@eecs.harvard.edu](mailto:vj@eecs.harvard.edu).

The challenge is the ecosystem’s many possible combinations of machine-learning tasks, models, data sets, frameworks, tool sets, libraries, architectures, and inference engines, which make inference benchmarking almost intractable. The spectrum of ML tasks is broad, including but not limited to image classification and localization, object detection and segmentation, machine translation, automatic speech recognition, text to speech, and recommendations. Even for a specific task, such as image classification, many ML models are viable. These models serve in a variety of scenarios that range from taking a single picture on a smartphone to continuously and concurrently detecting pedestrians through multiple cameras in an autonomous vehicle. Consequently, ML tasks have vastly different quality requirements and real-time-processing demands. Even implementations of functions and operations that the models typically rely on can be highly framework specific, and they increase the complexity of the design and the task.

Both academic and industrial organizations have developed ML inference benchmarks. Examples include AIMatrix (Alibaba, 2018), EEMBC MLMark (EEMBC, 2019), and AIX-PRTE (Principled Technologies, 2019) from industry, as well as AI Benchmark (Ignatov et al., 2019), TBD (Zhu et al., 2018), Fathom (Adolf et al., 2016), and DAWNbench (Coleman et al., 2017) from academia. Each one has made substantial contributions to ML benchmarking, but they were developed without input from ML-system designers. As a result, there is no consensus on representative models, metrics, tasks, and rules across these benchmarks. For example, some efforts focus too much on specific ML applications (e.g., computer vision) or specific domains (e.g., embedded inference). Moreover, it is important to devise the right performance metrics for inference so the evaluation accurately reflects how these models operate in practice. Latency, for instance, is the primary metric in many initial benchmarking efforts, but latency-bounded throughput is more relevant for many cloud inference scenarios.

Therefore, two critical needs remain unmet: (i) standard evaluation criteria for ML inference systems and (ii) an extensive (but reasonable) set of ML applications/models that cover existing inference systems across all major domains.

MLPerf Inference answers the call with a benchmark suite that complements MLPerf Training (Mattson et al., 2019). Jointly developed by the industry with input from academic researchers, more than 30 organizations as well as more than 200 ML engineers and practitioners assisted in the benchmark design and engineering process. This community architected MLPerf Inference to measure inference performance across a wide variety of ML hardware, software, systems, and services. The benchmark suite defines a set of tasks (models, data sets, scenarios, and quality targets) that represent real-world deployments, and it specifies the

evaluation metrics. In addition, the benchmark suite comes with permissive rules that allow comparison of different architectures under realistic scenarios.

Unlike traditional SPEC CPU-style benchmarks that run out of the box (Dixit, 1991), MLPerf promotes competition by allowing vendors to reimplement and optimize the benchmark for their system and then submit the results. To make results comparable, it defines detailed rules. It provides guidelines on how to benchmark inference systems, including when to start the performance-measurement timing, what preprocessing to perform before invoking the model, and which transformations and optimizations to employ. Such meticulous specifications help ensure comparability across ML systems because all follow the same rules.

We describe the design principles and architecture of the MLPerf Inference benchmark’s initial release (v0.5). We received over 600 submissions across a variety of tasks, frameworks, and platforms from 14 organizations. Audit tests validated the submissions, and the tests cleared 595 of them as valid. The final results show a four-orders-of-magnitude performance variation ranging from embedded devices and smartphones to data-center systems. MLPerf Inference adopts the following principles for a tailored approach to industry-standard benchmarking:

1. Pick representative workloads that everyone can access.
2. Evaluate systems in realistic scenarios.
3. Set target qualities and tail-latency bounds in accordance with real use cases.
4. Allow the benchmarks to flexibly showcase both hardware and software capabilities.
5. Permit the benchmarks to change rapidly in response to the evolving ML ecosystem.

The rest of the paper is organized as follows: Section 2 provides background, describing the differences in ML training versus ML inference and the challenges to creating a benchmark that covers the broad ML inference landscape. Section 3 describes the goals of MLPerf Inference. Section 4 presents MLPerf’s underlying inference-benchmark architecture and reveals the design choices for version 0.5. Section 5 summarizes the submission, review, and reporting process. Section 6 highlights v0.5 submission results to demonstrate that MLPerf Inference is a well-crafted industry benchmark. Section 7 shares the important lessons learned and prescribes a tentative roadmap for future work. Section 8 compares MLPerf Inference with prior efforts. Section 9 concludes the paper. Section 10 acknowledges the individuals who contributed to the benchmark’s development or validated the effort by submitting results.

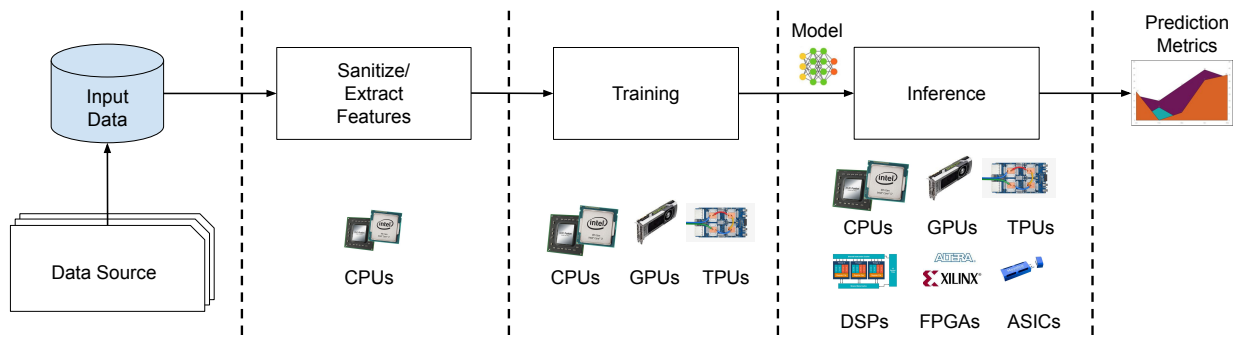


Figure 1. Stages of a typical ML pipeline. The first stage involves gathering data to train the models. The raw data is often noisy, so it requires processing before training a deep neural network (DNN). The hardware landscape for DNN training and inference is diverse.

## 2 BENCHMARKING CHALLENGES

We provide background on ML execution (Section 2.1) and describe the extreme heterogeneity that makes developing an ML inference benchmark challenging (Section 2.2).

### 2.1 ML Pipeline

Machine learning generally involves a series of complicated tasks (Figure 1). Nearly every ML pipeline begins by acquiring data to train and test the models. Raw data is typically sanitized and normalized before use because real-world data often contains errors, irrelevancies, or biases that reduce the quality and accuracy of ML models.

ML benchmarking focuses on two phases: training and inference. During training, models learn to make predictions from inputs. For example, a model may learn to predict the subject of a photograph or the most fluent translation of a sentence from English to German. During inference, models make predictions about their inputs, but they no longer learn. This phase is increasingly crucial as ML moves from research to practice, serving trillions of queries daily. Despite its apparent simplicity relative to training, the task of balancing latency, throughput, and accuracy for real-world applications makes optimizing inference difficult.

### 2.2 ML Inference Benchmarking Complexity

Creating a useful ML benchmark involves four critical challenges: (1) the diversity of models, (2) the variety of deployment scenarios, (3) the array of inference systems, and (4) the lack of a standard inference workflow.

#### 2.2.1 Diversity of Models

Even for a single task, such as image classification, numerous models present different trade-offs between accuracy and computational complexity, as Figure 2 shows. These

models vary tremendously in compute and memory requirements (e.g., a 50x difference in Gflops), while the corresponding Top-1 accuracy ranges from 55% to 83% (Bianco et al., 2018). This variation creates a Pareto frontier rather than one optimal choice.

Choosing the right model depends on the application. For example, pedestrian detection in autonomous vehicles has a much higher accuracy requirement than does labeling animals in photographs, owing to the different consequences of wrong predictions. Similarly, quality-of-service requirements for inference vary by several orders of magnitude from effectively no latency requirement for offline processes to milliseconds for real-time applications. Covering this design space necessitates careful selection of models that represent realistic scenarios.

Another challenge is that models vary wildly, so it is difficult to draw meaningful comparisons. In many cases, such as in Figure 2, a small accuracy change (e.g., a few percent) can drastically change the computational requirements (e.g., 5–10x). For example, SE-ResNeXt-50 (Hu et al., 2018; Xie et al., 2017) and Xception (Chollet, 2017) achieve roughly the same accuracy ( $\sim 79\%$ ) but exhibit a 2x difference in computational requirements ( $\sim 4$  Gflops versus  $\sim 8$  Gflops).

#### 2.2.2 Diversity of Deployment Scenarios

In addition to accuracy and computational complexity, the availability and arrival patterns of the input data vary with the deployment scenario. For example, in offline batch processing such as photo categorization, all the data may be readily available in (network) storage, allowing accelerators to reach and maintain peak performance. By contrast, translation, image tagging, and other web applications may experience variable arrival patterns based on end-user traffic.

Similarly, real-time applications such as augmented reality and autonomous vehicles handle a constant flow of data

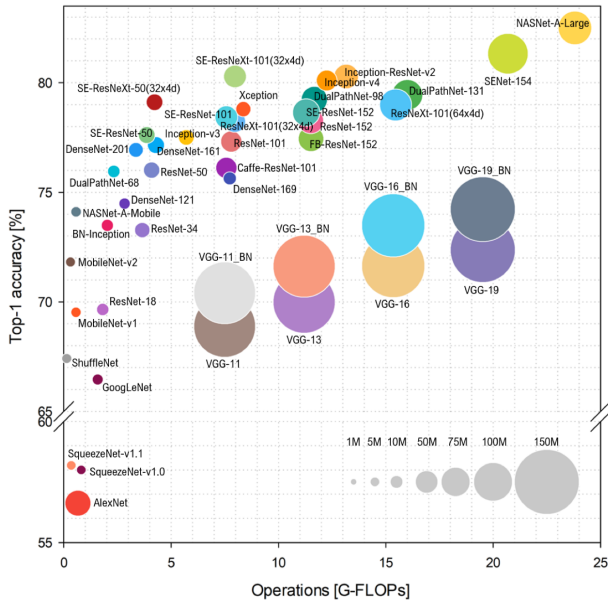


Figure 2. An example of ML-model diversity for image classification (figure from Bianco et al. (2018)). No single model is optimal; each one presents a unique design trade-off between accuracy, memory requirements, and computational complexity.

rather than having it all in memory. Although the same general model architecture could be employed in each scenario, data batching and similar optimizations may be inapplicable, leading to drastically different performance. Timing the on-device inference latency alone fails to reflect the real-world inference requirements.

2.2.3 Diversity of Inference Systems

The possible combinations of different inference applications, data sets, models, machine-learning frameworks, tool sets, libraries, systems, and platforms are numerous. Figure 3 shows the breadth and depth of the ML space. The hardware and software side exhibit substantial complexity.

On the software side, about a dozen ML frameworks commonly serve for developing deep-learning models, such as Caffe/Caffe2 (Jia et al., 2014), Chainer (Tokui et al., 2015), CNTK (Seide & Agarwal, 2016), Keras (Chollet et al., 2015), MXNet (Chen et al., 2015), TensorFlow (Abadi et al., 2016), and PyTorch (Paszke et al., 2017). Independently, there are also many optimized libraries, such as cuDNN (Chetlur et al., 2014), Intel MKL (Intel, 2018a), and FBGEMM (Khudia et al., 2018), supporting various inference run times, such as Apple CoreML (Apple, 2017), Intel OpenVino (Intel, 2018b), NVIDIA TensorRT (NVIDIA), ONNX Runtime (Bai et al., 2019), Qualcomm SNPE (Qualcomm), and TF-Lite (Lee et al., 2019a).

Each combination has idiosyncrasies that make supporting the most current neural-network model architectures a chal-

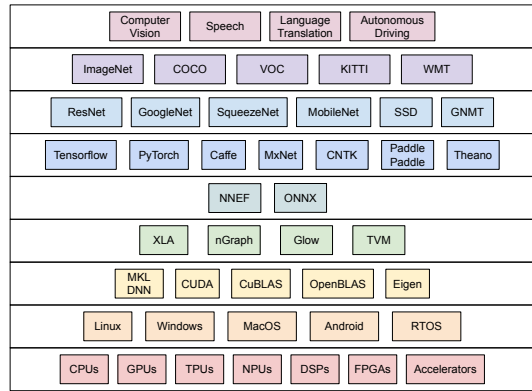


Figure 3. Software and hardware options at every level of the inference stack. The combinations across the layers make benchmarking ML inference systems a particularly challenging problem.

lenge. Consider the Non-Maximum Suppression (NMS) operator implementation for object detection. When training object-detection models in TensorFlow, the regular NMS operator smooths out imprecise bounding boxes for a single object. But this implementation is unavailable in TensorFlow Lite, which is tailored for mobile and instead implements fast NMS. As a result, when converting the model from TensorFlow to TensorFlow Lite, the accuracy of SSD-MobileNets-v1 decreases from 23.1% to 22.3% mAP. These types of subtle differences make it hard to port models exactly from one framework to another.

On the hardware side, platforms are tremendously diverse, ranging from familiar processors (e.g., CPUs, GPUs, and DSPs) to FPGAs, ASICs, and exotic accelerators such as analog and mixed-signal processors. Each platform comes with hardware-specific features and constraints that enable or disrupt performance depending on the model and scenario. Combining this diversity with the range of software systems above presents a unique challenge to deriving a robust and useful ML benchmark that meets industry needs.

2.2.4 Lack of a Standard Inference Workflow

There are many ways to optimize model performance. For example, quantizing floating-point weights decreases memory footprint and bandwidth requirements and increases computational throughput (wider vectors), but it also decreases model accuracy. Some platforms require quantization because they lack floating-point support. Low-power mobile devices, for example, call for such an optimization.

Other transformations are more complicated and change the network structure to boost performance further or exploit unique features of the inference platform. An example is reshaping image data from space to depth. The enormous variety of ML inference hardware and software means no one method can prepare trained models for all deployments.

### 3 MLPERF INFERENCE GOALS

To overcome the challenges, MLPerf Inference adopted a set of principles for developing a robust yet flexible benchmark suite based on community-driven development.

#### 3.1 Representative, Broadly Accessible Workloads

For the initial version 0.5, we chose tasks that reflect major commercial and research scenarios for a large class of submitters and that capture a broad set of computing motifs. To focus on the realistic rules and testing infrastructure, we selected a minimum-viable-benchmark approach to accelerate the development process. Where possible, we adopted models that were part of the MLPerf Training v0.6 suite (Mattson et al., 2019), thereby amortizing the benchmark-development effort.

The current version’s tasks and models are modest in scope. MLPerf Inference v0.5 comprises three tasks and five models: image classification (ResNet-50 (He et al., 2016) and MobileNet-v1 (Howard et al., 2017)), object detection (SSD-ResNet34—i.e., SSD (Liu et al., 2016) with a ResNet34 backbone—and SSD-MobileNet-v1—i.e., SSD with a MobileNet-v1 backbone), and machine translation (GNMT (Wu et al., 2016)). We plan to add others.

We chose our tasks and models through a consensus-driven process and considered community feedback to ensure their relevance. Our models are mature and have earned broad community support. Because the industry has studied them and can build efficient systems, benchmarking is accessible and provides a snapshot that shows the state of ML systems. Moreover, we focused heavily on the benchmark’s modular design to make adding new models and tasks less costly. As we show in Section 6.7, our design has allowed MLPerf Inference users to easily add new models. Our plan is to extend the scope to include more areas, tasks, models, and so on. Additionally, we aim to maintain consistency and alignment between the training and inference benchmarks.

#### 3.2 System Evaluation Using Realistic Scenarios

As our submission results show, ML inference systems vary in power consumption across four or more orders of magnitude and cover a wide variety of applications as well as physical deployments that range from deeply embedded devices to smartphones to data centers. The applications have a variety of usage models and many figures of merit, which in turn require multiple performance metrics. For example, the figure of merit for an image-recognition system that classifies a video camera’s output will be entirely different than for a cloud-based translation system. To address these various models, we surveyed MLPerf’s broad membership, which includes both customers and vendors. On the basis of that feedback, we identified four scenarios that represent

many critical inference applications.

Our goal is a method that simulates the realistic behavior of the inference system under test; such a feature is unique among AI benchmarks. To this end, we developed the Load Generator (LoadGen) tool, which is a query-traffic generator that mimics the behavior of real-world systems. It has four scenarios: *single-stream*, *multistream*, *server* and *offline*. They emulate the ML-workload behavior of mobile devices, autonomous vehicles, robotics, and cloud-based setups.

#### 3.3 Target Qualities and Tail-Latency Bounds

Quality and performance are intimately connected for all forms of machine learning, but the role of quality targets in inference is distinct from that in training. For training, the performance metric is the time to train to a specific quality, making accuracy a first-order consideration. For inference, the starting point is a pretrained reference model that achieves a target quality. Still, many system architectures can sacrifice model quality to achieve lower latency, lower total cost of ownership (TCO), or higher throughput.

The trade-offs between accuracy, latency, and TCO are application specific. Trading 1% model accuracy for 50% lower TCO is prudent when identifying cat photos, but it is less so during online pedestrian detection. For MLPerf, we define a model’s quality targets. To reflect this important aspect of real world-deployments, we established per-model and scenario targets for inference latency and model quality. The latency bounds and target qualities are based on input gathered from end users.

#### 3.4 Flexibility to Showcase Hardware and Software

Systems benchmarks can be characterized as language level (SPECInt (Dixit, 1991)), API level (LINPACK (Dongarra, 1988)), or semantic level (TPC (Council, 2005)). The ML community has embraced a wide variety of languages and libraries, so MLPerf Inference is a semantic-level benchmark. This type specifies the task to be accomplished and the general rules of the road, but it leaves implementation details to the submitters.

The MLPerf Inference benchmarks are flexible enough that submitters can optimize the reference models, run them through their preferred software tool chain, and execute them on their hardware of choice. Thus, MLPerf Inference has two divisions: closed and open. Strict rules govern the closed division, whereas the open division is more permissive and allows submitters to change the model, achieve different quality targets, and so on. The closed division is designed to address the lack of a standard inference-benchmarking workflow.

Within each division, submitters may file their results under specific categories on the basis of their hardware and

software components' availability. There are three system categories: available; preview; and research, development, or other systems. Systems in the first category are available off the shelf, while systems in the second category allow vendors to provide a sneak peek into their capabilities. At the other extreme are bleeding-edge ML solutions in the third category that are not ready for production use.

In summary, MLPerf Inference allows submitters to exhibit many different systems across varying product-innovation, maturity, and support levels.

### 3.5 Benchmark Changes for Rapidly Evolving ML

MLPerf Inference v0.5 is only the beginning. The benchmark will evolve. We are working to add more models (e.g., recommendation and time-series models), more scenarios (e.g., “burst” mode), better tools (e.g., a mobile application), and better metrics (e.g., timing preprocessing) to more accurately reflect the performance of the whole ML pipeline.

## 4 DESIGN AND IMPLEMENTATION

In this section we describe the design and implementation of MLPerf Inference v0.5. We also define the components of an inference system (Section 4.1) and detail how an inference query flows through one such system (Section 4.2). Our discussion also covers the MLPerf Inference tasks for v0.5 (Section 4.3).

### 4.1 Inference System Under Test (SUT)

A complete MLPerf Inference system contains multiple components: a data set, a system under test (SUT), the Load Generator (LoadGen), and an accuracy script. Figure 4 shows an overview of an MLPerf Inference system. The data set, LoadGen, and accuracy script are fixed for all submissions and are provided by MLPerf. Submitters have wide discretion to implement an SUT according to their architecture's requirements and their engineering judgment. By establishing a clear boundary between submitter-owned and MLPerf-owned components, the benchmark maintains comparability among submissions.

### 4.2 Life of a Query

At startup, the LoadGen requests that the SUT load samples into memory. The MLPerf Inference rules allow them to be loaded into DRAM as an untimed operation. The SUT loads the samples into DRAM and may perform other timed operations as the rules stipulate. These untimed operations may include but are not limited to compilation, cache warmup, and preprocessing.

The SUT signals the LoadGen when it is ready to receive the first query. A query is a request for inference on one

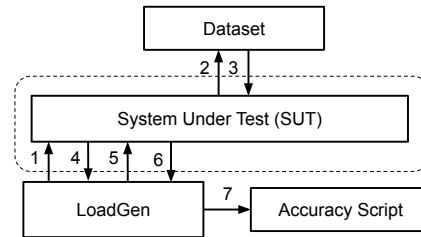


Figure 4. MLPerf Inference system under test (SUT) and how the components integrate. (1) The LoadGen requests that the SUT load samples; (2–3) the SUT loads samples into memory; (4) the SUT signals the LoadGen when it is ready; (5) the LoadGen issues requests to the SUT; (6) the benchmark processes the results and returns the results to the LoadGen; and (7) the LoadGen outputs logs, which the accuracy script then reads and verifies.

or more samples. The LoadGen sends queries to the SUT in accordance with the selected scenario. Depending on that scenario, it can submit queries one at a time, at regular intervals, or in a Poisson distribution.

The SUT runs inference on each query and sends the response back to the LoadGen, which either logs the response or discards it. After the run, an accuracy script checks the logged responses to determine whether the model accuracy is within tolerance.

We provide a clear interface between the SUT and LoadGen so new scenarios and experiments can be handled in the LoadGen and rolled out to all models and SUTs without extra effort. Doing so also facilitates compliance and auditing, since many technical rules about query arrivals, timing, and accuracy are implemented outside of submitter code. As we describe in Section 6.7, one submitter obtained results for over 60 image-classification and object-detection models.

Moreover, placing the performance-measurement code outside of submitter code is congruent with MLPerf's goal of end-to-end system benchmarking. To that end, the LoadGen measures the holistic performance of the entire SUT rather than any individual part. Finally, this condition enhances the benchmark's realism: inference engines typically serve as black-box components of larger systems.

### 4.3 Benchmark Tasks

Designing ML benchmarks is fundamentally different from designing non-ML benchmarks. MLPerf defines high-level tasks (e.g., image classification) that a machine-learning system can perform. For each one, we provide a canonical reference model in a few widely used frameworks. The reference model and weights offer concrete instantiations of the ML task, but formal mathematical equivalence is unnecessary. For example, a fully connected layer can be implemented with different cache-blocking and evaluation strategies. Consequently, submitting results requires opti-

## MLPerf Inference Benchmark

AREA	TASK	REFERENCE MODEL	DATA SET	QUALITY TARGET
VISION	IMAGE CLASSIFICATION (HEAVY)	RESNET-50 v1.5 25.6M PARAMETERS 7.8 GOPS / INPUT	IMAGENET (224x224)	99% OF FP32 (76.456%) TOP-1 ACCURACY
VISION	IMAGE CLASSIFICATION (LIGHT)	MOBILENET-V1 224 4.2M PARAMETERS 1.138 GOPS / INPUT	IMAGENET (224x224)	98% OF FP32 (71.676%) TOP-1 ACCURACY
VISION	OBJECT DETECTION (HEAVY)	SSD-RESNET34 36.3M PARAMETERS 433 GOPS / INPUT	COCO (1,200x1,200)	99% OF FP32 (0.20 MAP)
VISION	OBJECT DETECTION (LIGHT)	SSD-MOBILENET-V1 6.91M PARAMETERS 2.47 GOPS / INPUT	COCO (300x300)	99% OF FP32 (0.22 MAP)
LANGUAGE	MACHINE TRANSLATION	GNMT 210M PARAMETERS	WMT16 EN-DE	99% OF FP32 (23.9 SACREBLEU)

Table 1. ML Tasks in MLPerf Inference v0.5. Each one reflects critical commercial and research use cases for a large class of submitters, and together they also capture a broad set of computing motifs (e.g., CNNs and RNNs).

mizations to achieve good performance.

The concept of a reference model and a valid class of equivalent implementations creates freedom for most ML systems while still enabling relevant comparisons of inference systems. MLPerf provides reference models using 32-bit floating-point weights and, for convenience, also provides carefully implemented equivalent models to address the three most popular formats: TensorFlow (Abadi et al., 2016), PyTorch (Paszke et al., 2017), and ONNX (Bai et al., 2019).

As Table 1 illustrates, we selected a set of vision and language tasks along with associated reference models. We chose vision and translation because they are widely used across all computing systems, from edge devices to cloud data centers. Additionally, mature and well-behaved reference models with different architectures (e.g., CNNs and RNNs) were available.

For the vision tasks, we defined both heavyweight and lightweight models. The former are representative of systems with greater compute resources, such as a data center or autonomous vehicle, where increasing the computation cost for better accuracy is a reasonable trade-off. In contrast, the latter models are appropriate for systems with constrained compute resources and low latency requirements, such as smartphones and low-cost embedded devices.

For all tasks, we standardized on free and publicly available data sets to ensure the entire community can participate. Because of licensing restrictions on some data sets (e.g., ImageNet), we do not host them directly. Instead, the data is downloaded before running the benchmark.

### 4.3.1 Image Classification

Image classification is widely used in commercial applications and is also a de facto standard for evaluating ML-system performance. A classifier network takes an image as input and selects the class that best describes it. Example applications include photo searches, text extraction from images, and industrial automation, such as object sorting and defect detection.

For image classification, we use the standard ImageNet 2012 data set (Deng et al., 2009) and crop to 224x224 during preprocessing. We selected two models: a higher-accuracy and more computationally expensive heavyweight model as well as a computationally lightweight model that is faster but less accurate. Image-classification quality is the classifier’s Top-1 accuracy.

The heavyweight model, ResNet-50 v1.5 (He et al., 2016; MLPerf, 2019), comes directly from the MLPerf Training suite to maintain alignment. ResNet-50 is the most common network for performance claims. Unfortunately, it has multiple subtly different implementations that make most comparisons difficult. In our training suite, we specifically selected ResNet-50 v1.5 to ensure useful comparisons and compatibility across major frameworks. We also extensively studied and characterized the network for reproducibility and low run-to-run training variation, making it an obvious and low-risk choice.

The lightweight model, MobileNets-v1 224 (Howard et al., 2017), is built around smaller, depth-wise-separable convolutions to reduce the model complexity and computational burden. MobileNets is a family of models that offer

varying compute and accuracy options—we selected the full-width, full-resolution MobileNet-v1-1.0-224. This network reduces the parameters by 6.1x and the operations by 6.8x compared with ResNet-50 v1.5. We evaluated both MobileNet-v1 and v2 (Sandler et al., 2018) for the MLPerf Inference v0.5 suite and selected the former, as it has garnered wider adoption.

#### 4.3.2 Object Detection

Object detection is a complex vision task that determines the coordinates of bounding boxes around objects in an image and classifies the image. Object detectors typically use a pretrained image-classifier network as a backbone or a feature extractor, then perform regression for localization and bounding-box selection. Object detection is crucial for automotive applications, such as detecting hazards and analyzing traffic, and for mobile-retail tasks, such as identifying items in a picture.

For object detection, we chose the COCO data set (Lin et al., 2014) with both a lightweight and heavyweight model. Our small model uses the 300x300 image size, which is typical of resolutions in smartphones and other compact devices. For the larger model, we upscale the data set to more closely represent the output of a high-definition image sensor (1.44 MP total). The choice of the larger input size is based on community feedback, especially from automotive and industrial-automation customers. The quality metric for object detection is mean average precision (mAP).

The heavyweight object detector’s reference model is SSD (Liu et al., 2016) with a ResNet34 backbone, which also comes from our training benchmark. The lightweight object detector’s reference model uses a MobileNet-v1-1.0 backbone, which is more typical for constrained computing environments. We selected the MobileNet feature detector on the basis of feedback from the mobile and embedded communities.

#### 4.3.3 Translation

Neural machine translation (NMT) is popular in the rapidly evolving field of natural-language processing. NMT models translate a sequence of words from a source language to a target language and are used in translation applications and services. Our translation data set is WMT16 EN-DE (WMT, 2016). The quality measurement is Bilingual Evaluation Understudy Score (Bleu) (Papineni et al., 2002). In MLPerf Inference, we specifically employ SacreBleu (Post, 2018).

For the translation, we chose GNMT (Wu et al., 2016), which employs a well-established recurrent-neural-network (RNN) architecture and is part of the training benchmark. GNMT is representative of RNNs, which are popular for sequential and time-series data, and it ensures our reference-

model suite captures a wide variety of compute motifs.

## 4.4 Quality Targets

Many architectures can trade model quality for lower latency, lower TCO, or greater throughput. To reflect this important aspect of real-world deployments, we established per-model and scenario targets for latency and model quality. We adopted quality targets that for 8-bit quantization were achievable with considerable effort.

MLPerf Inference requires that almost all implementations achieve a quality target within 1% of the FP32 reference model’s accuracy (e.g., the ResNet-50 v1.5 model achieves 76.46% Top-1 accuracy, and an equivalent model must achieve at least 75.70% Top-1 accuracy). Initial experiments, however, showed that for mobile-focused networks—MobileNet and SSD-MobileNet—the accuracy loss was unacceptable without retraining. We were unable to proceed with the low accuracy because performance benchmarking would become unrepresentative.

To address the accuracy drop, we took three steps. First, we trained the MobileNet models for quantization-friendly weights, enabling us to narrow the quality window to 2%. Second, to reduce the training sensitivity of MobileNet-based submissions, we provided equivalent MobileNet and SSD-MobileNet implementations quantized to an 8-bit integer format. Third, for SSD-MobileNet, we reduced the quality requirement to 22.0 mAP to account for the challenges of using MobileNets as a backbone.

To improve the submission comparability, we disallow retraining. Our prior experience and feasibility studies confirmed that for 8-bit integer arithmetic, which was an expected deployment path for many systems, the  $\sim 1\%$  relative-accuracy target was easily achievable without retraining.

## 4.5 Scenarios and Metrics

The diverse inference applications have various usage models and figures of merit, which in turn require multiple performance metrics. To address these models, we specify four scenarios that represent important inference applications. Each one has a unique performance metric, as Table 2 illustrates. The LoadGen discussed in Section 4.7 simulates the scenarios and measures the performance.

**Single-stream.** This scenario represents one inference-query stream with a query sample size of one, reflecting the many client applications where responsiveness is critical. An example is offline voice transcription on Google’s Pixel 4 smartphone. To measure performance, the LoadGen injects a single query; when the query is complete, it records the completion time and injects the next query. The performance metric is the query stream’s 90th-percentile latency.

## MLPerf Inference Benchmark

SCENARIO	QUERY GENERATION	METRIC	SAMPLES/QUERY	EXAMPLES
SINGLE-STREAM (SS)	SEQUENTIAL	90TH-PERCENTILE LATENCY	1	TYPING AUTOCOMPLETE, REAL-TIME AR
MULTISTREAM (MS)	ARRIVAL INTERVAL WITH DROPPING	NUMBER OF STREAMS SUBJECT TO LATENCY BOUND	$N$	MULTICAMERA DRIVER ASSISTANCE, LARGE-SCALE AUTOMATION
SERVER (S)	POISSON DISTRIBUTION	QUERIES PER SECOND SUBJECT TO LATENCY BOUND	1	TRANSLATION WEBSITE
OFFLINE (O)	BATCH	THROUGHPUT	AT LEAST 24,576	PHOTO CATEGORIZATION

Table 2. Scenario description and metrics. Each scenario targets a real-world use case based on customer and vendor input.

**Multistream.** This scenario represents applications with a stream of queries, but each query comprises multiple inferences, reflecting a variety of industrial-automation and remote-sensing applications. For example, many autonomous vehicles analyze frames from six to eight cameras that stream simultaneously.

To model a concurrent scenario, the LoadGen sends a new query comprising  $N$  input samples at a fixed time interval (e.g., 50 ms). The interval is benchmark specific and also acts as a latency bound that ranges from 50 to 100 milliseconds. If the system is available, it processes the incoming query. If it is still processing the prior query in an interval, it skips the interval and delays the remaining queries by one interval.

No more than 1% of the queries may produce one or more skipped intervals. A query’s  $N$  input samples are contiguous in memory, which accurately reflects production input pipelines and avoids penalizing systems that would otherwise require that samples be copied to a contiguous memory region before starting inference. The performance metric is the integer number of streams that the system supports while meeting the QoS requirement.

**Server.** This scenario represents online server applications where query arrival is random and latency is important. Almost every consumer-facing website is a good example, including services such as online translation from Baidu, Google, and Microsoft. For this scenario, the load generator sends queries, with one sample each, in accordance with a Poisson distribution. The SUT responds to each query within a benchmark-specific latency bound that varies from 15 to 250 milliseconds. No more than 1% of queries may exceed the latency bound for the vision tasks and no more than 3% may do so for translation. The server scenario’s performance metric is the Poisson parameter that indicates the queries per second achievable while meeting the QoS requirement.

**Offline.** This scenario represents batch-processing applications where all the input data is immediately available and latency is unconstrained. An example is identifying the

people and locations in a photo album. For the offline scenario, the LoadGen sends to the system a single query that includes all sample-data IDs to be processed, and the system is free to process the input data in any order. Similar to the multistream scenario, neighboring samples in the query are contiguous in memory. The metric for the offline scenario is throughput measured in samples per second.

For the multistream and server scenarios, latency is a critical component of the system behavior and will constrain various performance optimizations. For example, most inference systems require a minimum (and architecture-specific) batch size to achieve full utilization of the underlying computational resources. But in a server scenario, the arrival rate of inference queries is random, so systems must carefully optimize for tail latency and potentially process inferences with a suboptimal batch size.

Table 3 shows the relevant latency constraints for each task in v0.5. As with other aspects of MLPerf, we selected these constraints on the basis of community consultation and feasibility assessments. The multistream arrival times for most vision tasks correspond to a frame rate of 15–20 Hz, which is a minimum for many applications. The server QoS constraints derive from estimates of the inference timing budget given an overall user latency target.

TASK	MULTISTREAM ARRIVAL TIME	SERVER QoS CONSTRAINT
IMAGE CLASSIFICATION (HEAVY)	50 MS	15 MS
IMAGE CLASSIFICATION (LIGHT)	50 MS	10 MS
OBJECT DETECTION (HEAVY)	66 MS	100 MS
OBJECT DETECTION (LIGHT)	50 MS	10 MS
MACHINE TRANSLATION	100 MS	250 MS

Table 3. Latency constraints for each task in the multistream and server scenarios.

TAIL-LATENCY PERCENTILE	CONFIDENCE INTERVAL	ERROR MARGIN	INFERENCES	ROUNDED INFERENCES
90%	99%	0.50%	23,886	$3 \times 2^{13} = 24,576$
95%	99%	0.25%	50,425	$7 \times 2^{13} = 57,344$
99%	99%	0.05%	262,742	$33 \times 2^{13} = 270,336$

Table 4. Query requirements for statistical confidence. All results must meet the minimum LoadGen scenario requirements.

#### 4.6 Statistical Confidence

To ensure our results are statistically robust and adequately capture steady-state system behavior, each task and scenario combination requires a minimum number of queries. That number is determined by the tail-latency percentile, the desired margin, and the desired confidence interval.

Confidence is the probability that a latency bound is within a particular margin of the reported result. We chose a 99% confidence bound and set the margin to a value much less than the difference between the tail-latency percentage and 100%. Conceptually, that margin ought to be relatively small. Thus, we selected a margin that is one-twentieth of the difference between the tail-latency percentage and 100%.

The equation is as follows:

$$Margin = \frac{1 - TailLatency}{20} \quad (1)$$

$$NumQueries = (Normslnv(\frac{1 - Confidence}{2}))^2 \times \frac{TailLatency \times (TailLatency - 1)}{Margin^2} \quad (2)$$

Table 4 shows the query requirements. The total query count and tail-latency percentile are scenario and task specific. The single-stream scenario only requires 1,024 queries, and the offline scenario requires a single query containing at least 24,576 samples. The single-stream scenario has the fewest queries to execute because we wanted the run time to be short enough that embedded platforms and smartphones could complete the runs quickly.

For scenarios with latency constraints, our goal is to ensure a 99% confidence interval that the constraints hold. As a result, the benchmarks with more-stringent latency constraints require more queries in a highly nonlinear fashion. The number of queries is based on the aforementioned statistics and is rounded up to the nearest multiple of  $2^{13}$ .

A 99<sup>th</sup>-percentile guarantee requires 262,742 queries, which rounds up to  $33 \times 2^{13}$ , or 270K. For both multistream and server, this guarantee for vision tasks requires 270K queries, as Table 5 shows. Because a multistream benchmark will process  $N$  samples per query, the total number of samples

MODEL	NUMBER OF QUERIES / SAMPLES PER QUERY			
	SINGLE-STREAM	MULTISTREAM	SERVER	OFFLINE
IMAGE CLASSIFICATION (HEAVY)	1K / 1	270K / $N$	270K / 1	1 / 24K
IMAGE CLASSIFICATION (LIGHT)	1K / 1	270K / $N$	270K / 1	1 / 24K
OBJECT DETECTION (HEAVY)	1K / 1	270K / $N$	270K / 1	1 / 24K
OBJECT DETECTION (LIGHT)	1K / 1	270K / $N$	270K / 1	1 / 24K
MACHINE TRANSLATION	1K / 1	90K / $N$	90K / 1	1 / 24K

Table 5. Number of queries and samples per query for each task.

will be  $N \times 270K$ . Machine translation has a 97<sup>th</sup>-percentile latency guarantee and requires only 90K queries.

For repeatability, we run both the multistream and server scenarios several times. But the multistream scenario’s arrival rate and query count guarantee a 2.5- to 7-hour run time. To strike a balance between repeatability and run time, we require five runs for the server scenario, with the result being the minimum of these five runs. The other scenarios require one run. We expect to revisit this choice in future benchmark versions.

All benchmarks must also run for at least 60 seconds and process additional queries and/or samples as the scenarios require. The minimum run time ensures they will measure the equilibrium behavior of power-management systems and systems that support dynamic voltage and frequency scaling (DVFS), particularly for the single-stream scenario with a small number of queries.

#### 4.7 Load Generator

The LoadGen is a traffic generator that loads the SUT and measures performance. Its behavior is controlled by a configuration file it reads at the start of the benchmark run. The LoadGen produces the query traffic according to the rules of the previously described scenarios (i.e., single-stream, multistream, server, and offline). Additionally, the LoadGen collects information for logging, debugging, and postprocessing the data. It records queries and responses from the SUT, and at the end of the run, it reports statistics, summarizes the results, and determines whether the run was valid.

Figure 5 shows how the LoadGen generates query traffic for each scenario. In the server scenario, for instance, it issues queries in accordance with a Poisson distribution to mimic a server’s query-arrival rates. In the single-stream case, it issues a query to the SUT and waits for completion of that query before issuing another.

##### 4.7.1 Design

MLPerf will evolve, introducing new tasks and removing old ones as the field progresses. Accordingly, the LoadGen’s design is flexible enough to handle changes to the inference-

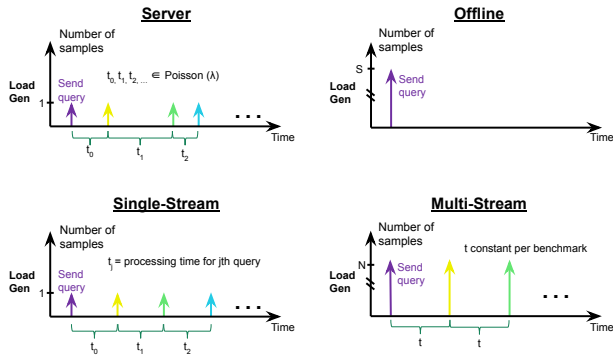


Figure 5. The timing and number of queries from the Load Generator (LoadGen) vary between benchmark scenarios. All five ML tasks can run in any one of the four scenarios.

task suite. We achieve this feat by decoupling the LoadGen from the benchmarks and the internal representations (e.g., the model, scenarios, and quality and latency metrics).

The LoadGen is implemented as a standalone C++ module with well-defined APIs; the benchmark calls it through these APIs (and vice versa through callbacks). This decoupling at the API level allows it to easily support various language bindings, permitting benchmark implementations in any language. Presently, the LoadGen supports Python, C, and C++ bindings; additional bindings can be added.

Another major benefit of decoupling the LoadGen from the benchmark is that the LoadGen is extensible to support more scenarios. Currently, MLPerf supports four of them; we may add more, such as a multitенancy mode where the SUT must continuously serve multiple models while maintaining QoS constraints.

#### 4.7.2 Implementation

The LoadGen abstracts the details of the data set (e.g., images) behind sample IDs. Data-set samples receive an index between 0 and  $N$ . A query represents the smallest input unit that the benchmark ingests from the LoadGen. It consists of one or more data-set sample IDs, each with a corresponding response ID to differentiate between multiple instances of the same sample.

The rationale for a response ID is that for any given task and scenario—say, an image-classification multistream scenario—the LoadGen may reissue the same data (i.e., an image with a unique sample ID) multiple times across the different streams. To differentiate between them, the LoadGen must assign different reference IDs to accurately track when each sample finished processing.

At the start, the LoadGen directs the benchmark to load a list of samples into memory. Loading is untimed and the SUT may also perform allowed data preprocessing. The

LoadGen then issues queries, passing sample IDs to the benchmark for execution on the inference hardware. The queries are pre-generated to reduce overhead during the timed portion of the test.

As the benchmark finishes processing the queries, it informs the LoadGen through a function named *QuerySamplesComplete*. The LoadGen makes no assumptions regarding how the SUT may partition its work, so any thread can call this function with any set of samples in any order. *QuerySamplesComplete* is thread safe, is wait-free bounded, and makes no syscalls, allowing it to scale recording to millions of samples per second and to minimize the performance variance introduced by the LoadGen, which would affect long-tail latency.

The LoadGen maintains a logging thread that gathers events as they stream in from other threads. At the end of the benchmark run, it outputs a set of logs that report the performance and accuracy stats.

#### 4.7.3 Operating Modes

The LoadGen has two primary operating modes: accuracy and performance. Both are necessary to make a valid MLPerf submission.

**Accuracy mode.** The LoadGen goes through the entire data set for the ML task. The model’s task is to run inference on the complete data set. Afterward, accuracy results appear in the log files, ensuring that the model met the required quality target.

**Performance mode.** The LoadGen avoids going through the entire data set, as the system’s performance can be determined by subjecting it to enough data-set samples.

#### 4.7.4 Validation Features

The LoadGen has features that ensure the submission system complies with the rules. In addition, it can self-check to determine whether its source code has been modified during the submission process. To facilitate validation, the submitter provides an experimental config file that allows use of non-default LoadGen features. For v0.5, the LoadGen enables the following four tests.

**Accuracy verification.** The purpose of this test is to ensure valid inferences in performance mode. By default, the results that the inference system returns to the LoadGen are not logged and thus are not checked for accuracy. This choice reduces or eliminates processing overhead to allow accurate measurement of the inference system’s performance. In this test, results returned from the SUT to the LoadGen are logged randomly. The log is checked against the log generated in accuracy mode to ensure consistency.

**On-the-fly caching detection.** By default, LoadGen pro-

duces queries by randomly selecting with replacement from the data set, and inference systems may receive queries with duplicate samples. This outcome is likely for high-performance systems that process many samples relative to the data-set size. To represent realistic deployments, the MLPerf rules prohibit caching of queries or intermediate data. The test has two parts. The first part generates queries with unique sample indices. The second generates queries with duplicate sample indices. Performance is measured in each case. The way to detect caching is to determine whether the test with duplicate sample indices runs significantly faster than the test with unique sample indices.

**Alternate-random-seed testing.** In ordinary operation, the LoadGen produces queries on the basis of a fixed random seed. Optimizations based on that seed are prohibited. The alternate-random-seed test replaces the official random seed with alternates and measures the resulting performance.

## 4.8 Model Equivalence

The goal of MLPerf Inference is to measure realistic system-level performance across a wide variety of architectures. But the four properties of realism, comparability, architecture neutrality, and friendliness to small submission teams require careful trade-offs.

Some inference deployments involve teams of compiler, computer-architecture, and machine-learning experts aggressively co-optimizing the training and inference systems to achieve cost, accuracy, and latency targets across a massive global customer base. An unconstrained inference benchmark, however, would disadvantage companies with less experience and fewer ML-training resources.

Therefore, we set the model-equivalence rules to allow submitters to, for efficiency, reimplement models on different architectures. The rules provide a complete list of disallowed techniques and a list of allowed technique examples. We chose an explicit blacklist to encourage a wide range of techniques and to support architectural diversity. The list of examples illustrates the boundaries of the blacklist while also encouraging common and appropriate optimizations.

Examples of allowed techniques include the following: arbitrary data arrangement as well as different input and in-memory representations of weights, mathematically equivalent transformations (e.g., tanh versus logistic, ReluX versus ReluY, and any linear transformation of an activation function), approximations (e.g., replacing a transcendental function with a polynomial), processing queries out of order within the scenario’s limits, replacing dense operations with mathematically equivalent sparse operations, fusing or unfusing operations, dynamically switching between one or more batch sizes, mixing experts that combine differently quantized weights.

### 4.8.1 Prohibited Optimizations

MLPerf Inference currently prohibits retraining and pruning to ensure comparability, although this restriction may fail to reflect realistic deployment for some large companies. The interlocking requirements to use reference weights (possibly with calibration) and minimum accuracy targets are most important for ensuring comparability in the closed division. The open division explicitly allows retraining and pruning.

We prohibit caching to simplify the benchmark design. In practice, real inference systems cache queries. For example, “*I love you*” is one of Google Translate’s most frequent queries, but the service does not translate the phrase *ab initio* each time. Realistically modeling caching in a benchmark, however, is a challenge because cache hit rates vary substantially with the application. Furthermore, our data sets are relatively small, and large systems could easily cache them in their entirety.

We also prohibit optimizations that are benchmark aware or data-set aware and that are inapplicable to production environments. For example, real query traffic is unpredictable, but for the benchmark, the traffic pattern is predetermined by the pseudorandom-number-generator seed. Optimizations that take advantage of a fixed number of queries or that use knowledge of the LoadGen implementation are prohibited. Similarly, any optimization employing statistical knowledge of the performance or accuracy data sets is prohibited. Finally, we disallow any technique that takes advantage of the upscaled images in the 1,200x1,200 COCO data set for the heavyweight object detector.

### 4.8.2 Preprocessing and Data Types

Ideally, a whole-system benchmark should capture all performance-relevant operations. MLPerf, however, explicitly allows untimed preprocessing. There is no vendor- or application-neutral preprocessing. For example, systems with integrated cameras can use hardware/software co-design to ensure that images arrive in memory in an ideal format; systems accepting JPEGs from the Internet cannot.

In the interest of architecture and application neutrality, we adopted a permissive approach to untimed preprocessing. Implementations may transform their inputs into system-specific ideal forms as an untimed operation.

MLPerf explicitly allows and enables quantization to a wide variety of numerical formats to ensure architecture neutrality. Submitters must pre-register their numerics to help guide accuracy-target discussions. The approved list for the closed division includes INT4, INT8, INT16, UINT8, UINT16, FP11 (sign, 5-bit mantissa, and 5-bit exponent), FP16, bfloat16, and FP32.

Quantization to lower-precision formats typically requires

calibration to ensure sufficient inference quality. For each reference model, MLPerf provides a small, fixed data set that can be used to calibrate a quantized network. Additionally, it offers MobileNet versions that are prequantized to INT8, since without retraining (which we disallow) the accuracy falls dramatically.

## 5 SUBMISSION, REVIEW, AND REPORTING

In this section, we describe the submission process for MLPerf Inference v0.5 (Sections 5.1). All submissions are peer reviewed for validity (Section 5.2). Finally, we describe how we report the results to the public (Section 5.3).

### 5.1 Submission

An MLPerf Inference submission contains information about the SUT: performance scores, benchmark code, a system-description file that highlights the SUT’s main configuration characteristics (e.g., accelerator count, CPU count, software release, and memory system), and LoadGen log files detailing the performance and accuracy runs for a set of task and scenario combinations. All this data is uploaded to a public GitHub repository for peer review and validation before release.

MLPerf Inference is a suite of tasks and scenarios that ensures broad coverage, but a submission can contain subset tasks and scenarios. Many traditional benchmarks, such as SPEC CPU, require submissions for all their components. This approach is logical for a general-purpose processor that runs arbitrary code, but ML systems are often highly specialized. For example, some are solely designed for vision or wake-word detection and cannot run other network types. Others target particular scenarios, such as a single-stream application, and are not intended for server-style applications (or vice versa). Accordingly, we allow submitters flexibility in selecting tasks and scenarios.

#### 5.1.1 Divisions

MLPerf Inference has two divisions for submitting results: closed and open. Submitters can send results to either or both, but they must use the same data set. The open division, however, allows free model selection and unrestricted optimization to foster ML-system innovation.

**Closed division.** The closed division enables comparisons of different systems. Submitters employ the same models, data sets, and quality targets to ensure comparability across wildly different architectures. This division requires preprocessing, postprocessing, and a model that is equivalent to the reference implementation. It also permits calibration for quantization (using the calibration data set we provide) and prohibits retraining.

**Open division.** The open division fosters innovation in ML systems, algorithms, optimization, and hardware/software co-design. Submitters must still perform the same ML task, but they may change the model architecture and the quality targets. This division allows arbitrary pre- and postprocessing and arbitrary models, including techniques such as retraining. In general, submissions are not directly comparable with each other or with closed submissions. Each open submission must include documentation about how it deviates from the closed division. *Caveat emptor!*

#### 5.1.2 Categories

Submitters must classify their submissions into one of three categories on the basis of hardware- and software-component availability: available; preview; and research, development, or other systems. This requirement helps consumers of the results identify the systems’ maturity level and whether they are readily available (either for rent online or for purchase).

**Available systems.** Available systems are generally the most mature and have stringent hardware- and software-availability requirements.

An available *cloud system* must have accessible pricing (either publicly or by request), have been rented by at least one third party, have public evidence of availability (e.g., a web page or company statement saying the product is available), and be “reasonably available” for additional third parties to rent by the submission date.

An *on-premise system* is available if all its components that substantially determine ML performance are available either individually or in aggregate (development boards that meet the substantially determined clause are allowed). An available component or system must have available pricing (either publicly advertised or available by request), have been shipped to at least one third party, have public evidence of availability (e.g., a web page or company statement saying the product is available), and be “reasonably available” for purchase by additional third parties by the submission date. In addition, submissions for on-premises systems must describe the system and its components in sufficient detail so that third parties can build a similar system.

Available systems must use a publicly available software stack consisting of the software components that substantially determine ML performance but are absent from the source code. An available software component must be well supported for general use and available for download.

**Preview systems.** Preview systems contain components that will meet the criteria for the available category within 180 days or by the next submission cycle, whichever is later. This restriction applies to both the hardware and software requirements. The goal of the preview category is to en-

able participants to submit results for new systems without burdening product-development cycles with the MLPerf schedule. Any system submitted to preview must then be submitted to available during the next cycle.

**Research, development, or other systems.** Research, development, or other (RDO) systems contain components not intended for production or general availability. An example is a prototype system that is a proof of concept. An RDO system includes one or more RDO components. These components submitted in one cycle may not be submitted as available until the third cycle or until 181 days have passed, whichever is later.

## 5.2 Review and Validation

MLPerf Inference submissions are self- and peer-reviewed for compliance with all rules. Compliance issues are tracked and raised with submitters, who must resolve them and then resubmit results.

A challenge of benchmarking inference systems is that many include proprietary and closed-source components, such as inference engines and quantization flows, that make peer review difficult. To accommodate these systems while ensuring reproducible results that are free from common errors, we developed a validation suite to assist with peer review.

Our validation tools perform experiments that help determine whether a submission complies with the defined rules. MLPerf Inference provides a suite of validation tests that submitters must run to qualify their submission as valid. MLPerf v0.5 tests the submission system using LoadGen validation features (Section 4.7.4).

In addition to LoadGen’s validation features, we use custom data sets to detect result caching. This behavior is validated by replacing the reference data set with a custom data set. We measure the quality and performance of the system operating on this custom data set and compare the results with operation on the reference data set.

## 5.3 Reporting

All results are published on the MLPerf website following review and validation. MLPerf Inference does not require that submitters include results for all the ML tasks. Therefore, some systems lack results for certain tasks and scenarios.

MLPerf Inference does not provide a “summary score.” Often in benchmarking, there is a strong desire to distill the capabilities of a complex system to a single score to enable a comparison of different systems. But not all ML tasks are equally important for all systems, and the job of weighting some more heavily than others is highly subjective.

At best, weighting and summarization are driven by the submitter catering to unique customer needs, as some systems

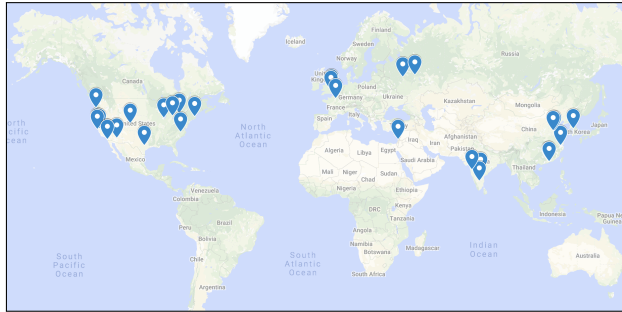


Figure 6. MLPerf Inference’s accessibility and global reach. The organizations responding to the v0.5 call for submissions hail from around the world, including the United States, Canada, the European Union, Russia, the Middle East, India, China, and South Korea. This domestic and international adoption reflects the community’s perspective that the benchmark is comprehensive and scientifically rigorous, and worthy of engineering time for submissions.

may be optimized for specific ML tasks. For instance, some real-world systems are more highly optimized for vision than for translation. In such scenarios, averaging the results across all tasks makes no sense, as the submitter may not be targeting particular markets.

## 6 RESULTS

We received over 600 submissions in all three categories (available, preview, and RDO) across the closed and open divisions. Our results are the most extensive corpus of inference performance data available to the public, covering a range of ML tasks and scenarios, hardware architectures, and software run times. Each has gone through extensive review before receiving approval as a valid MLPerf result. After review, we cleared 595 results as valid.

We evaluated the closed-division results on the basis of four of the five objectives our benchmark aimed to achieve. The exception is setting target qualities and tail-latency bounds in accordance with real use cases, which we do not discuss because a static benchmark setting applies to every inference task. Omitting that isolated objective, we present our analysis as follows:

- Pick representative workloads that everyone can access (Sections 6.1 and 6.2).
- Evaluate systems in realistic scenarios (Section 6.3).
- Allow the benchmark to flexibly showcase both hardware and software capabilities (Sections 6.4, 6.5, and 6.6).
- Permit the benchmark to change rapidly in response to the evolving ML ecosystem (Section 6.7).

## 6.1 Accessibility and Global Reach

A primary goal for MLPerf Inference was to create a widely available benchmark. To this end, the first round of submissions came from 14 worldwide organizations, hailing from the United States, Canada, Russia, the European Union, the Middle East, India, China, and South Korea, as Figure 6 shows.

The submitters represent many organizations that range from startups to original equipment manufacturers (OEMs), cloud-service providers, and system integrators. They include Alibaba, Centaur Technology, Dell EMC, dividiti, FuriosaAI, Google, Habana, Hailo, Inspur, Intel, NVIDIA, Polytechnic University of Milan, Qualcomm, and Tencent.

## 6.2 Task Coverage

MLPerf Inference v0.5 submitters are allowed to pick any task to evaluate their system’s performance. The distribution of results across tasks can thus reveal whether those tasks are of interest to ML-system vendors.

We analyzed the submissions to determine the overall task coverage. Figure 7 shows the breakdown for the tasks and models in the closed division. Although the most popular model was, unsurprisingly, ResNet-50 v1.5, it was just under three times as popular as GNMT, the least popular model. This small spread and the otherwise uniform distribution suggests we selected a representative set of tasks.

In addition to selecting representative tasks, another goal is to provide vendors with varying quality and performance targets. Depending on the use case, the ideal ML model may differ (as Figure 2 shows, a vast range of models can target a given task). Our results reveal that vendors equally supported different models for the same task because each model has unique quality and performance trade-offs. In the case of object detection, we saw the same number of submissions for both SSD-MobileNet-v1 and SSD-ResNet34.

## 6.3 Scenario Usage

We aim to evaluate systems in realistic use cases—a major motivator for the LoadGen (Section 4.7) and scenarios (see Section 4.5). To this end, Table 6 shows the distribution of results across the various task and scenario combinations.

Across all the tasks, the single-stream and offline scenarios are the most widely used and are also the easiest to optimize and run. Server and multistream were more complicated and had longer run times because of the QoS requirements and more-numerous queries.

GNMT garnered no multistream submissions, possibly because the constant arrival interval is unrealistic in machine translation. Therefore, it was the only model and scenario

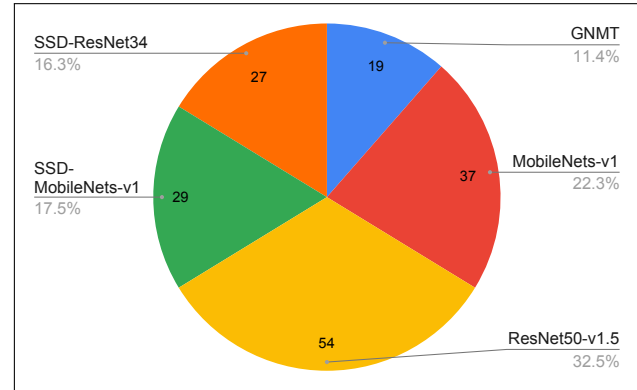


Figure 7. Results from the closed division. The distribution of models indicates MLPerf Inference capably selected representative workloads for the initial v0.5 benchmark release.

combination with no submissions.

## 6.4 Processor Types

Machine-learning solutions can be deployed on a variety of platforms, ranging from fully general-purpose CPUs to programmable GPUs and DSPs, FPGAs, and fixed-function accelerators. Our results reflect this diversity.

Figure 8 shows that the MLPerf Inference submissions covered most hardware categories. The system diversity indicates that our inference benchmark suite and method for v0.5 can evaluate any processor architecture.

## 6.5 Software Frameworks

In addition to the various hardware types are many ML software frameworks. Table 7 shows the variety of frameworks used to benchmark the hardware platforms. ML software plays a vital role in unleashing the hardware’s performance.

Some run times are specifically designed to work with certain types of hardware to fully harness their capabilities; employing the hardware without the corresponding framework may still succeed, but the performance may fall short of the hardware’s potential. The table shows that CPUs have

	SINGLE-STREAM	MULTISTREAM	SERVER	OFFLINE
GNMT	2	0	6	11
MOBILENET-V1	18	3	5	11
RESNET-50 V1.5	19	5	10	20
SSD-MOBILENET-V1	8	3	5	13
SSD-RESNET34	4	4	7	12
TOTAL	51	15	33	67

Table 6. Closed-division submissions for the tasks and LoadGen scenarios. The high coverage of models and scenarios implies that the benchmark captures important real-world use cases.

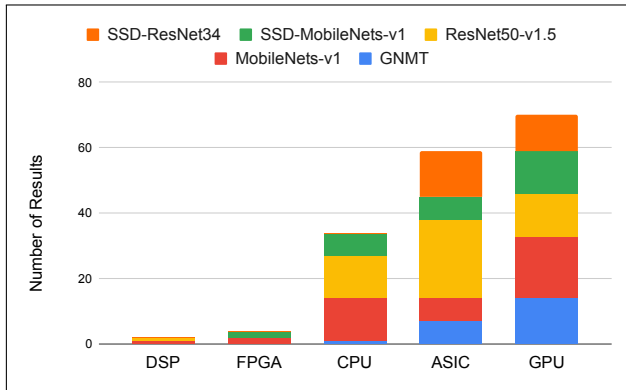


Figure 8. Results from the closed division. The results cover many processor architectures. Almost every kind—CPUs, GPUs, DSPs, FPGAs, and ASICs—appeared in the submissions.

the most framework diversity and that TensorFlow has the most architectural variety.

## 6.6 Diversity of Systems

The MLPerf Inference v0.5 submissions cover a broad range of systems on the power and performance scale, from mobile and edge devices to cloud computing. The performance delta between the smallest and largest inference systems is four orders of magnitude, or about 10,000x.

Table 8 shows the performance range for each task and scenario in the closed division (except for GNMT, which had no multistream submissions). For example, in the case of ResNet-50 v1.5 offline, the highest-performing system is over 10,000x faster than the lowest-performing one. Unsurprisingly, the former comprised multiple ML accelerators, whereas the latter was a low-power laptop-class CPU. This delta for single-stream is surprising given that additional accelerators cannot reduce latency, and it reflects an even more extensive range of systems than the other scenarios. In particular, the single-stream scenario includes many smartphone processors, which target very low power.

Figure 9 shows the results across all tasks and scenarios. In cases such as the MobileNet-v1 single-stream scenario (SS), ResNet-50 v1.5 SS, and SSD-MobileNet-v1 SS, systems exhibit a large performance difference (100x). Because these models have many applications, the systems that target them cover everything from low-power embedded devices to high-performance servers. GNMT server (S) shows much less performance variation between systems.

The broad performance range implies that the selected tasks (as a starting point) for MLPerf Inference v0.5 are general enough to represent a variety of use cases and market segments. The wide array of systems also indicates that our method (LoadGen, metrics, etc.) is broadly applicable.

	ASIC	CPU	DSP	FPGA	GPU
ARM NN		X			X
FURIOSA-AI				X	
HAILO SDK	X				
HANGUANG-AI	X				
ONNX		X			
OPENVINO		X			
PYTORCH		X			
SNPE			X		
SYNAPSE	X				
TENSORFLOW	X	X			X
TF-LITE		X			
TENSORRT					X

Table 7. Summary of software framework versus hardware architecture in the closed division. The hardware benchmarking involves many different frameworks. Preventing submitters from reimplementing the benchmark would have made it impossible to support the diversity of systems tested.

## 6.7 Open Division

The open division is the vanguard of MLPerf’s benchmarking efforts. It is less rigid than the closed division; we received over 400 results. The submitters ranged from startups to large organizations.

A few highlights from the open division are the use of 4-bit quantization to boost performance, an exploration of a wide range of models to perform the ML task (instead of using the reference model), and a demonstration of one system’s ability to deliver high throughput even under tighter latency bounds—tighter than those in the closed-division rules.

In addition, we received a submission that pushed the limits of mobile-chipset performance. Typically, most vendors use one accelerator at a time to do inference. In this case, a vendor concurrently employed multiple accelerators to deliver high throughput in a multistream scenario—a rarity in conventional mobile use cases. Nevertheless, it shows that the MLPerf Inference open division is encouraging the

	SINGLE-STREAM	MULTISTREAM	SERVER	OFFLINE
GNMT	2	N/A	5	2,367
MOBILENET-V1	1,199	29	9	438
RESNET-50V1.5	11154	27	26	10,289
SSD-MOBILENET-V1	8	36	25	657
SSD-RESNET34	8	44	9	147

Table 8. Closed-division performance summary across tasks and scenarios. Each entry is the ratio of the highest to lowest performance. The performance range is as much as 10,000x. GNMT appears as N/A for multistream because it had no submissions.

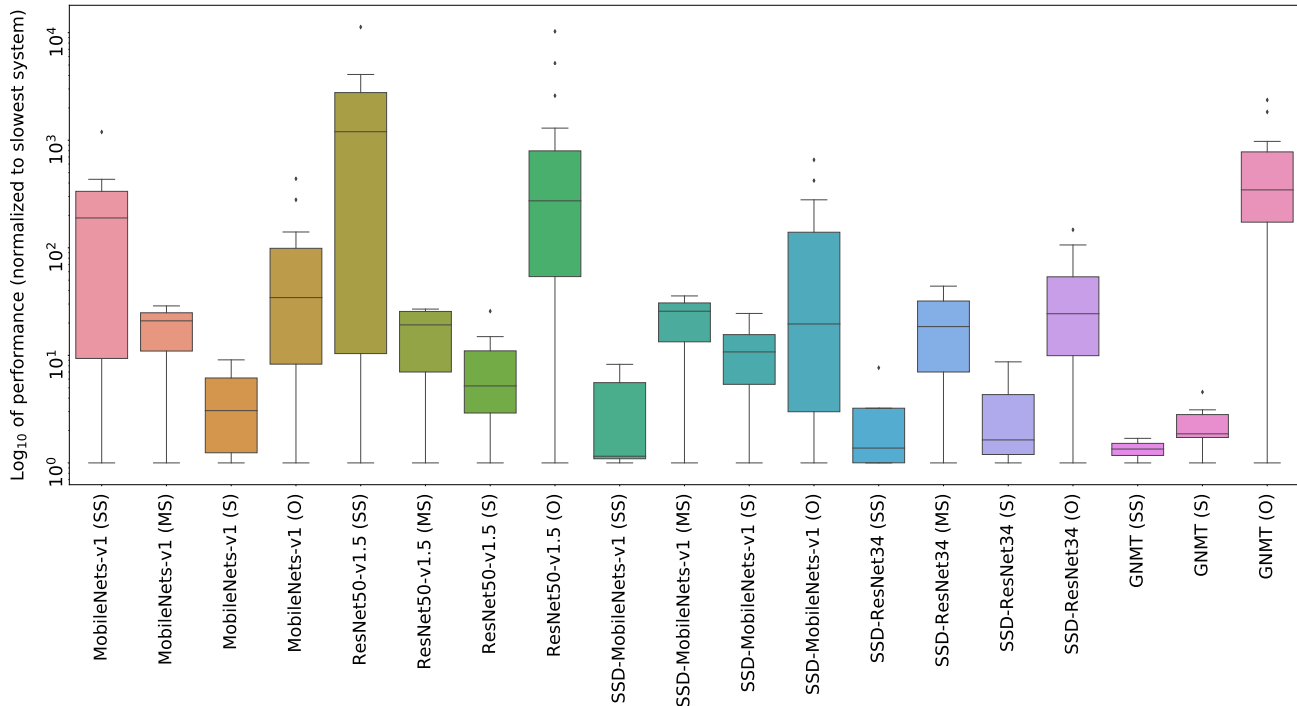


Figure 9. Results from the closed division. Normalized performance distribution on log scale ( $\log_{10}$ ) across models for the single-stream (SS), multistream (MS), server (S), and offline (O) scenarios. The boxplot shows the performance distribution of all system submissions for a specific model and scenario combination. The results are normalized to the slowest system representing that combination. A wide range emerges across all tasks and scenarios. GNMT MS is absent because no submitter ran the multistream scenario.

industry to push the limits of systems.

In yet another interesting submission, two organizations jointly evaluated 12 object-detection models—YOLO v3 (Redmon & Farhadi, 2018), Faster-RCNN (Ren et al., 2015) with a variety of backbones, and SSD (Liu et al., 2016) with a variety of backbones—on a desktop platform. The open-division results save practitioners and researchers from having to manually perform similar explorations, while also showcasing potential techniques and optimizations.

## 7 LESSONS LEARNED

We reflect on our v0.5 benchmark-development effort and share some lessons we learned from the experience.

### 7.1 Community-Driven Benchmark Development

There are two main approaches to building an industry-standard benchmark. One is to create the benchmark in house, release it, and encourage the community to adopt it. The other is first to consult the community and then build the benchmark through a consensus-based effort. The former approach is useful when seeding an idea, but the latter is necessary to develop an industry-standard benchmark. MLPerf Inference employed the latter.

MLPerf Inference began as a community-driven effort on July 12, 2018. We consulted more than 15 organizations. Since then, many other organizations have joined the MLPerf Inference working group. Applying the wisdom of several ML engineers and practitioners, we built the benchmark from the ground up, soliciting input from the ML-systems community as well as hardware end users. This collaborative effort led us to directly address the industry’s diverse needs from the start. For instance, the LoadGen and scenarios emerged from our desire to span the many inference-benchmark needs of various organizations.

Although convincing competing organizations to agree on a benchmark is a challenge, it is still possible—as MLPerf Inference shows. Every organization has unique requirements and expectations, so reaching a consensus was sometimes tricky. In the interest of progress, everyone agreed to make decisions on the basis of “grudging consensus.” These decisions were not always in favor of any one organization. Organizations would comply to keep the process moving or defer their requirements to a future version so benchmark development could continue.

Ultimately, MLPerf Inference exists because competing organizations saw beyond their self-interest and worked together to achieve a common goal: establishing the best

ways to measure ML inference performance.

## 7.2 Perfect Is the Enemy of Good

MLPerf Inference v0.5 has a modest number of tasks and models. Early in the development process, it was slated to cover 11 ML tasks: image classification, object detection, speech recognition, machine translation, recommendation, text (e.g., sentiment) classification, language modeling, text to speech, face identification, image segmentation, and image enhancement. We chose these tasks to cover the full breadth of ML applications relevant to the industry.

As it matured, however, engineering hurdles and the participating organizations’ benchmark-carrying capacity limited our effort. The engineering hurdles included specifying and developing the LoadGen system, defining the scenarios, and building the reference implementations. The LoadGen, for instance, involved 11 engineers from nine organizations. The reference implementations involved 34 people from 15 organizations contributing to our GitHub repository.

We deemed that overcoming the engineering hurdles was a priority, as they would otherwise limit our ability to represent various workloads and to grow in the long term. Hence, rather than incorporating many tasks and models right away, we trimmed the number of tasks to five and focused on developing a proper method and infrastructure.

With the hurdles out of the way, a small team or even an individual can add new models. For instance, thanks to the LoadGen and a complementary workflow-automation technology (Fursin et al., 2016), one MLPerf contributor with only three employees swept more than 60 computer-vision models in the open division.

Similarly, adding another task would require only a modest effort to integrate with the LoadGen and implement the model. This flexibility allows us to accommodate the changing ML landscape, and it saves practitioners and researchers from having to perform these explorations manually, all while showcasing potential techniques and optimizations for future versions of the closed division.

## 7.3 Audits and Auditability

MLPerf is committed to integrity through rigorous submitter cross-auditing and to the privacy of the auditing process. This process was uncontentious and smooth flowing. Three innovations helped ease the audit process: permissive rules, the LoadGen, and the submission checker.

Concerns arose during rule-making that submitters would discover loopholes in the blacklist, allowing them to “break” the benchmark and, consequently, undermine the legitimacy of the entire MLPerf project. Submitters worked together to patch loopholes as they appeared because all are invested in

the success of the benchmark.

The LoadGen improved auditability by separating measurement and experimental setup into a shared component. The only possible error in the experimental procedure is use of the wrong LoadGen settings. The LoadGen, therefore, significantly reduced compliance issues.

Finally, MLPerf provided a script for checking submissions. The script allowed submitters to verify that they submitted all required files in the right formats along with the correct directory layouts. It also verified LoadGen settings and scanned logs for noncompliance.

The submission-checker script kept all submissions relatively uniform and allowed submitters to quickly identify and resolve potential problems. In future revisions, MLPerf will aim to expand the range of issues the submission script discovers. We also plan to include additional checker scripts and tools to further smooth the audit process.

## 8 PRIOR ART IN AI/ML BENCHMARKING

The following summary describes prior AI/ML inference benchmarking. Each of these benchmarks has made unique contributions. MLPerf has strived to incorporate and build on the best aspects of previous work while ensuring it includes community input. Compared with earlier efforts, MLPerf brings more-rigorous performance metrics that we carefully selected for each major use case along with a much wider (but still compact) set of ML applications and models based on the community’s input.

**AI Benchmark.** AI Benchmark (Ignatov et al., 2019) is arguably the first mobile-inference benchmark suite. It covers 21 computer-vision and AI tests grouped in 11 sections. These tests are predominantly computer-vision tasks (image recognition, face detection, and object detection), which are also well represented in the MLPerf suite. The AI Benchmark results and leaderboard focus primarily on Android smartphones and only measure inference latency. The suite provides a summary score, but it does not explicitly specify the quality targets. Relative to AI Benchmark, we aim at a wider variety of devices (submissions for v0.5 range from IoT devices to server-scale systems) and multiple scenarios. Another important distinction is that MLPerf does not endorse a summary score, as we mentioned previously.

**EEMBC MLMark.** EEMBC MLMark (EEMBC, 2019) is an ML benchmark suite designed to measure the performance and accuracy of embedded inference devices. It includes image-classification (ResNet-50 v1 and MobileNet-v1) and object-detection (SSD-MobileNet-v1) workloads, and its metrics are latency and throughput. Its latency and throughput modes are roughly analogous to the MLPerf single-stream and offline modes. MLMark measures per-

formance at explicit batch sizes, whereas MLPerf allows submitters to choose the best batch sizes for different scenarios. Also, the former imposes no target-quality restrictions, whereas the latter imposes stringent restrictions.

**Fathom.** An early ML benchmark, Fathom (Adolf et al., 2016) provides a suite of neural-network models that incorporate several types of layers (e.g., convolution, fully connected, and RNN). Still, it focuses on throughput rather than accuracy. Fathom was an inspiration for MLPerf: in particular, we likewise included a suite of models that comprise various layer types. Compared with Fathom, MLPerf provides both PyTorch and TensorFlow reference implementations for optimization, ensuring that the models in both frameworks are equivalent, and it also introduces a variety of inference scenarios with different performance metrics.

**AIXPRT.** Developed by Principled Technologies, AIXPRT (Principled Technologies, 2019) is a closed, proprietary AI benchmark that emphasizes ease of use. It consists of image-classification, object-detection, and recommender workloads. AIXPRT publishes prebuilt binaries that employ specific inference frameworks on supported platforms. The goal of this approach is apparently to allow technical press and enthusiasts to quickly run the benchmark. Binaries are built using Intel OpenVino, TensorFlow, and NVIDIA TensorRT tool kits for the vision workloads, as well as MXNet for the recommendation system. AIXPRT runs these workloads using FP32 and INT8 numbers with optional batching and multi-instance, and it evaluates performance by measuring latency and throughput. The documentation and quality requirements are unpublished but are available to members. In contrast, MLPerf tasks are supported on any framework, tool kit, or OS; they have precise quality requirements; and they work with a variety of scenarios.

**AI Matrix.** AI Matrix (Alibaba, 2018) is Alibaba’s AI-accelerator benchmark for both cloud and edge deployment. It takes the novel approach of offering four benchmark types. First, it includes micro-benchmarks that cover basic operators such as matrix multiplication and convolutions that come primarily from DeepBench. Second, it measures performance for common layers, such as fully connected layers. Third, it includes numerous full models that closely track internal applications. Fourth, it offers a synthetic benchmark designed to match the characteristics of real workloads. The full AI Matrix models primarily target TensorFlow and Caffe, which Alibaba employs extensively and which are mostly open source. We have a smaller model collection and focus on simulating scenarios using LoadGen.

**DeepBench.** Microbenchmarks such as DeepBench (Baidu, 2017) measure the library implementation of kernel-level operations (e.g., 5,124x700x2,048 GEMM) that are important for performance in production models. They are useful for efficient model development but fail to address the com-

plexity of testing and evaluating full ML models.

**TBD (Training Benchmarks for DNNs).** TBD (Zhu et al., 2018) is a joint project of the University of Toronto and Microsoft Research that focuses on ML training. It provides a wide spectrum of ML models in three frameworks (TensorFlow, MXNet, and CNTK), along with a powerful tool chain for their improvement. It primarily focuses on evaluating GPU performance and only has one full model (Deep Speech 2) that covers inference. We considered including TBD’s Deep Speech 2 model but lacked the time.

**DawnBench.** DawnBench (Coleman et al., 2017) was the first multi-entrant benchmark competition to measure the end-to-end performance of deep-learning systems. It allowed optimizations across model architectures, optimization procedures, software frameworks, and hardware platforms. DawnBench inspired MLPerf, but our benchmark offers more tasks, models, and scenarios.

To summarize, MLPerf Inference builds on the best of prior work and improves on it, in part through community-driven feedback (Section 7.1). The result has been new features, such as the LoadGen (which can run models in different scenarios), the open and closed divisions, and so on.

## 9 CONCLUSION

More than 200 ML researchers, practitioners, and engineers from academia and industry helped to bring the MLPerf Inference benchmark from concept (June 2018) to result submission (October 2019). This team, drawn from 32 organizations, developed the reference implementations and rules, and submitted over 600 performance measurements gathered on a wide range of systems. Of these performance measurements, 595 cleared the audit process as valid submissions and were approved for public consumption.

MLPerf Inference v0.5 is just the beginning. The key to any benchmark’s success, especially in a rapidly changing field such as ML, is a development process that can respond quickly to changes in the ecosystem. Work has already started on the next version. We expect to update the current models (e.g., MobileNet-v1 to v2), expand the list of tasks (e.g., recommendation), increase the processing requirements by scaling the data-set sizes (e.g., 2 MP for SSD large), allow aggressive performance optimizations (e.g., retraining for quantization), simplify benchmarking through better infrastructure (e.g., a mobile app), and increase the challenge to systems by improving the metrics (e.g., measuring power and adjusting the quality targets).

We welcome your input and contributions. Visit the MLPerf website (<https://mlperf.org>) for additional details. Results for v0.5 are available online ([https://github.com/mlperf/inference\\_results\\_v0.5](https://github.com/mlperf/inference_results_v0.5)).

**10 ACKNOWLEDGEMENTS**

MLPerf Inference is the work of many individuals from multiple organizations. In this section, we acknowledge all those who helped produce the first set of results or supported the overall benchmark development.

**ALIBABA T-HEAD**

Zhi Cai, Danny Chen, Liang Han, Jimmy He, David Mao, Benjamin Shen, ZhongWei Yao, Kelly Yin, XiaoTao Zai, Xiaohui Zhao, Jesse Zhou, and Guocai Zhu.

**BAIDU**

Newsha Ardalani, Ken Church, and Joel Hestness.

**CADENCE**

Debajyoti Pal.

**CENTAUR TECHNOLOGY**

Bryce Arden, Glenn Henry, CJ Holthaus, Kimble Houck, Kyle O'Brien, Parviz Palangpour, Benjamin Seroussi, and Tyler Walker.

**DELL EMC**

Frank Han, Bhavesh Patel, Vilmara Rocio Sanchez, and Rengan Xu.

**DIVIDITI**

Grigori Fursin and Leo Gordon.

**FACEBOOK**

Soumith Chintala, Kim Hazelwood, Bill Jia, and Sean Lee.

**FURIOSAAI**

Dongsun Kim and Sol Kim.

**GOOGLE**

Michael Banfield, Victor Bittorf, Bo Chen, Dehao Chen, Ke Chen, Chiachen Chou, Sajid Dalvi, Suyog Gupta, Blake Hechtman, Terry Heo, Andrew Howard, Sachin Joglekar, Allan Knies, Naveen Kumar, Cindy Liu, Thai Nguyen, Tayo Oguntebi, Yuechao Pan, Mangpo Phothilimthana, Jue Wang, Shibo Wang, Tao Wang, Qiumin Xu, Cliff Young, Ce Zheng, and Zongwei Zhou.

**HAILO**

Ohad Agami, Mark Grobman, and Tamir Tapuhi.

**INTEL**

Md Faijul Amin, Thomas Atta-fosu, Haim Barad, Barak Battash, Amit Bleiweiss, Maor Busidan, Deepak R Canchi, Baishali Chaudhuri, Xi Chen, Elad Cohen, Xu Deng, Pradeep Dubey, Matthew Eckelman, Alex Fradkin, Daniel Franch, Srujana Gattupalli, Xiaogang Gu, Amit Gur, MingXiao Huang, Barak Hurwitz, Ramesh Jaladi, Rohit Kalidindi, Lior Kalman, Manasa Kankanala, Andrey Karpenko, Noam Korem, Evgeny Lazarev, Hongzhen Liu, Guokai Ma, Andrey Malyshev, Manu Prasad Manmanthan, Ekaterina Mastrosova, Jerome Mitchell, Arijit Mukhopadhyay, Jitender Patil, Reuven Richman, Rachitha Prem Seelin, Maxim Shevtshov, Avi Shimalkovski, Dan Shirron, Hui Wu, Yong Wu, Ethan Xie, Cong Xu, Feng Yuan, and Eliran Zimmerman.

**MEDIATEK**

Bing Yu.

**MICROSOFT**

Scott McKay, Tracy Sharpe, and Changming Sun.

**MYRTLE**

Peter Baldwin.

**NVIDIA**

Felix Abecassis, Vikram Anjur, Jeremy Appleyard, Julie Bernauer, Anandi Bharwani, Ritika Borkar, Lee Bushen, Charles Chen, Ethan Cheng, Melissa Collins, Niall Emmart, Michael Fertig, Prashant Gaikwad, Anirban Ghosh, Mitch Harwell, Po-Han Huang, Wenting Jiang, Patrick Judd, Prethvi Kashinkunti, Milind Kulkarni, Garvit Kulshreshta, Jonas Li, Allen Liu, Kai Ma, Alan Menezes, Maxim Milakov, Rick Napier, Brian Nguyen, Ryan Olson, Robert Overman, Jhalak Patel, Brian Pharris, Yujia Qi, Randall Radmer, Supriya Rao, Scott Ricketts, Nuno Santos, Madhumita Sridhara, Markus Tavenrath, Rishi Thakka, Ani Vaidya, KS Venkatraman, Jin Wang, Chris Wilkerson, Eric Work, and Bruce Zhan.

**POLITECNICO DI MILANO**

Emanuele Vitali.

**QUALCOMM**

Srinivasa Chaitanya Gopireddy, Pradeep Jilagam, Chirag Patel, Harris Teague, and Mike Tremaine.

**SAMSUNG**

Rama Harihara, Jungwook Hong, David Tannenbaum, Simon Waters, and Andy White.

**STANFORD UNIVERSITY**

Peter Bailis and Matei Zaharia.

**SUPERMICRO**

Srini Bala, Ravi Chintala, Alec Duroy, Raju Penumatcha, Gayatri Pichai, and Sivanagaraju Yarramaneni.

**UNAFFILIATED**

Michael Gschwind and Justin Sang.

**UNIVERSITY OF CALIFORNIA, BERKELEY / GOOGLE**

David Patterson.

**XILINX**

Ziheng Gao, Yiming Hu, Satya Keerthi Chand Kudupudi, Ji Lu, Lu Tian, and Treeman Zheng.

**REFERENCES**

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, volume 16, pp. 265–283, 2016.

Adolf, R., Rama, S., Reagen, B., Wei, G.-Y., and Brooks, D. Fathom: Reference Workloads for Modern Deep Learning Methods. In *Workload Characterization (IISWC), 2016 IEEE International Symposium on*, pp. 1–10. IEEE, 2016.

Alibaba. Ai matrix. <https://aimatrix.ai/en-us/>, 2018.

Amodei, D. and Hernandez, D. Ai and compute. <https://blog.openai.com/ai-and-compute/>, 2018.

Apple. Core ml: Integrate machine learning models into your app. <https://developer.apple.com/documentation/coreml>, 2017.

Badrinarayanan, V., Kendall, A., and Cipolla, R. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.

Bai, J., Lu, F., Zhang, K., et al. Onnx: Open neural network exchange. <https://github.com/onnx/onnx>, 2019.

Baidu. DeepBench: Benchmarking Deep Learning Operations on Different Hardware. <https://github.com/baidu-research/DeepBench>, 2017.

Bianco, S., Cadene, R., Celona, L., and Napoletano, P. Benchmark analysis of representative deep neural network architectures. *IEEE Access*, 6:64270–64277, 2018.

Chen, T., Li, M., Li, Y., Lin, M., Wang, N., Wang, M., Xiao, T., Xu, B., Zhang, C., and Zhang, Z. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*, 2015.

Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., and Shelhamer, E. cudnn: Efficient primitives for deep learning. *CoRR*, abs/1410.0759, 2014. URL <http://arxiv.org/abs/1410.0759>.

Chollet, F. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.

Chollet, F. et al. Keras. <https://keras.io>, 2015.

- Coleman, C., Narayanan, D., Kang, D., Zhao, T., Zhang, J., Nardi, L., Bailis, P., Olukotun, K., Ré, C., and Zaharia, M. DAWNbench: An End-to-End Deep Learning Benchmark and Competition. *NIPS ML Systems Workshop*, 2017.
- Council, T. P. P. Transaction processing performance council. *Web Site*, <http://www.tpc.org>, 2005.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dixit, K. M. The spec benchmarks. *Parallel computing*, 17 (10-11):1195–1209, 1991.
- Dongarra, J. The linpack benchmark: An explanation. In *Proceedings of the 1st International Conference on Supercomputing*, pp. 456–474, London, UK, UK, 1988. Springer-Verlag. ISBN 3-540-18991-2. URL <http://dl.acm.org/citation.cfm?id=647970.742568>.
- EEMBC. Introducing the eembc mlmark benchmark. <https://www.eembc.org/mlmark/index.php>, 2019.
- Fursin, G., Lokhmotov, A., and Plowman, E. Collective knowledge: towards r&d sustainability. In *2016 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 864–869. IEEE, 2016.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- Han, S., Liu, X., Mao, H., Pu, J., Pedram, A., Horowitz, M. A., and Dally, W. J. Eie: efficient inference engine on compressed deep neural network. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pp. 243–254. IEEE, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Hu, J., Shen, L., and Sun, G. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.
- Ignatov, A., Timofte, R., Kulik, A., Yang, S., Wang, K., Baum, F., Wu, M., Xu, L., and Van Gool, L. Ai benchmark: All about deep learning on smartphones in 2019. *arXiv preprint arXiv:1910.06663*, 2019.
- Intel. Intel math kernel library. <https://software.intel.com/en-us/mkl>, 2018a.
- Intel. Intel distribution of openvino toolkit. <https://software.intel.com/en-us/openvino-toolkit>, 2018b.
- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T. Caffe: Convolutional Architecture for Fast Feature Embedding. In *ACM International Conference on Multimedia*, pp. 675–678. ACM, 2014.
- Khudia, D. S., Basu, P., and Deng, S. Open-sourcing fbgemm for state-of-the-art server-side inference. <https://engineering.fb.com/ml-applications/fbgemm/>, 2018.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lee, J., Chirkov, N., Ignasheva, E., Pisarchyk, Y., Shieh, M., Riccardi, F., Sarokin, R., Kulik, A., and Grundmann, M. On-device neural net inference with mobile gpus. *arXiv preprint arXiv:1907.01989*, 2019a.
- Lee, K., Rao, V., and Arnold, W. C. Accelerating facebook’s infrastructure with application-specific hardware. <https://engineering.fb.com/data-center-engineering/accelerating-infrastructure/>, 3 2019b.
- Levine, S., Pastor, P., Krizhevsky, A., Ibarz, J., and Quillen, D. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 37(4-5):421–436, 2018.
- Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.

- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 2014.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., and Berg, A. C. Ssd: Single shot multibox detector. In *European conference on computer vision*, pp. 21–37. Springer, 2016.
- Mattson, P., Cheng, C., Coleman, C., Diamos, G., Micikevicius, P., Patterson, D., Tang, H., Wei, G.-Y., Bailis, P., Bittorf, V., Brooks, D., Chen, D., Dutta, D., Gupta, U., Hazelwood, K., Hock, A., Huang, X., Jia, B., Kang, D., Kanter, D., Kumar, N., Liao, J., Narayanan, D., Oguntebi, T., Pekhimenko, G., Pentecost, L., Reddi, V. J., Robie, T., John, T. S., Wu, C.-J., Xu, L., Young, C., and Zaharia, M. MLperf training benchmark, 2019.
- MLPerf. MLPerf Reference: ResNet in TensorFlow. [https://github.com/mlperf/training/tree/master/image\\_classification/tensorflow/official](https://github.com/mlperf/training/tree/master/image_classification/tensorflow/official), 2019.
- Molchanov, P., Tyree, S., Karras, T., Aila, T., and Kautz, J. Pruning convolutional neural networks for resource efficient inference. *arXiv preprint arXiv:1611.06440*, 2016.
- NVIDIA. Nvidia tensorrt: Programmable inference accelerator. <https://developer.nvidia.com/tensorrt>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pp. 311–318. Association for Computational Linguistics, 2002.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.
- Post, M. A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771*, 2018.
- Principled Technologies. Aixprt community preview. <https://www.principledtechnologies.com/benchmarkxpert/aixprt/>, 2019.
- Qualcomm. Snapdragon neural processing engine sdk reference guide. <https://developer.qualcomm.com/docs/snpe/overview.html>.
- Redmon, J. and Farhadi, A. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- Ren, S., He, K., Girshick, R., and Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pp. 91–99, 2015.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018.
- Seide, F. and Agarwal, A. Cntk: Microsoft’s open-source deep-learning toolkit. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2135–2135. ACM, 2016.
- Tokui, S., Oono, K., Hido, S., and Clayton, J. Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems (LearningSys) in the twenty-ninth annual conference on neural information processing systems (NIPS)*, volume 5, pp. 1–6, 2015.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- WMT. First conference on machine translation, 2016. URL <http://www.statmt.org/wmt16/>.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.
- Xu, D., Anguelov, D., and Jain, A. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 244–253, 2018.
- Zhu, H., Akrouf, M., Zheng, B., Pelegrin, A., Jayarajan, A., Phanishayee, A., Schroeder, B., and Pekhimenko, G. Benchmarking and analyzing deep neural network training. In *2018 IEEE International Symposium on Workload Characterization (IISWC)*, pp. 88–100. IEEE, 2018.