



XR BENCH: AN EXTENDED REALITY (XR) MACHINE LEARNING BENCHMARK SUITE FOR THE METAVERSE

Hyoukjun Kwon¹ Krishnakumar Nair² Jamin Seo^{3*} Jason Yik^{4*}
Debabrata Mohapatra² Dongyuan Zhan² Jinook Song² Peter Capak² Peizhao Zhang²
Peter Vajda² Colby Banbury⁴ Mark Mazumder⁴ Liangzhen Lai² Ashish Sirasao²
Tushar Krishna³ Harshit Khaitan² Vikas Chandra² Vijay Janapa Reddi⁴

ABSTRACT

Real-time multi-task multi-model (MTMM) workloads, a new form of deep learning inference workloads, are emerging for applications areas like extended reality (XR) to support metaverse use cases. These workloads combine user interactivity with computationally complex machine learning (ML) activities. Compared to standard ML applications, these ML workloads present unique difficulties and constraints. Real-time MTMM workloads impose heterogeneity and concurrency requirements on future ML systems and devices, necessitating the development of new capabilities. This paper begins with a discussion of the various characteristics of these real-time MTMM ML workloads and presents an ontology for evaluating the performance of future ML hardware for XR systems. Next, we present XR BENCH, a collection of MTMM ML tasks, models, and usage scenarios that execute these models in three representative ways: cascaded, concurrent, and cascaded-concurrent for XR use cases. Finally, we emphasize the need for new metrics that capture the requirements properly. We hope that our work will stimulate research and lead to the development of a new generation of ML systems for XR use cases. XR BENCH is available as an open-source project: <https://github.com/XRBench>

1 INTRODUCTION

Applications based on machine learning (ML) are becoming prevalent. The number of ML models that must be supported on the edge, mobile, and data centers is growing. The success of ML across tasks in vision and speech recognition is furthering the development of increasingly sophisticated use cases. For instance, the *metaverse* (Meta, 2022c) combines multiple unit use cases (e.g., image classification and speech recognition) to create more sophisticated use cases such as real-time interactivity via virtual reality. Such sophisticated use cases demand more functionality, for which application engineers are increasingly relying on composability; rather than developing different large models for use cases, they are combining multiple smaller and specialized ML models to compose task functionality (Barham et al., 2022).

In this paper, we focus on this new class of ML workloads referred to as multi-task multi-model (MTMM) ML workloads, specifically in the context of extended reality (XR)

for metaverse use cases. A real-time MTMM application for extended reality is illustrated in Figure 1. The figure depicts how several MTMM models can be cascaded and operated concurrently, sometimes dynamically subject to certain conditions, to provide complex application-level functionality. The center section of the figure demonstrates that processing throughput requirements can vary depending on the usage scenario. The right side of the figure shows how there can be a variety of interleaved execution patterns for each of the concurrent jobs. MTMM workloads exhibit model heterogeneity, expanded computation scheduling spaces (Kwon et al., 2021), and usage-dependent real-time constraints, which makes them challenging to support compared to today’s single-task single-model (STSM) workloads.

We identify three key issues that arise with MTMM workloads that present interesting system-level design challenges. The first is *scenario-driven behavior*. All ML pipelines operate at a set frames per second (FPS) processing rate that is determined by a particular use case (e.g., virtual reality gaming, augmented reality social interaction, and outdoor activity recognition). A scenario may sometimes even demand zero FPS (i.e., deactivating a model) for models not required for the scenario. This fluctuating FPS is due to the context-based behavior that drives system resource utilization, which presents a hurdle when designing the underlying DNN accelerator—the heterogeneous workload makes it difficult to employ traditional DNN specialization.

*Equal contribution ¹EECS, University of California, Irvine, Irvine, California, USA ²Meta, Menlo Park, California, USA ³ECE, Georgia Institute of Technology, Atlanta, Georgia, USA ⁴SEAS, Harvard University, Cambridge, Massachusetts, USA. Correspondence to: Hyoukjun Kwon <hyoukjun.kwon.p@gmail.com>.

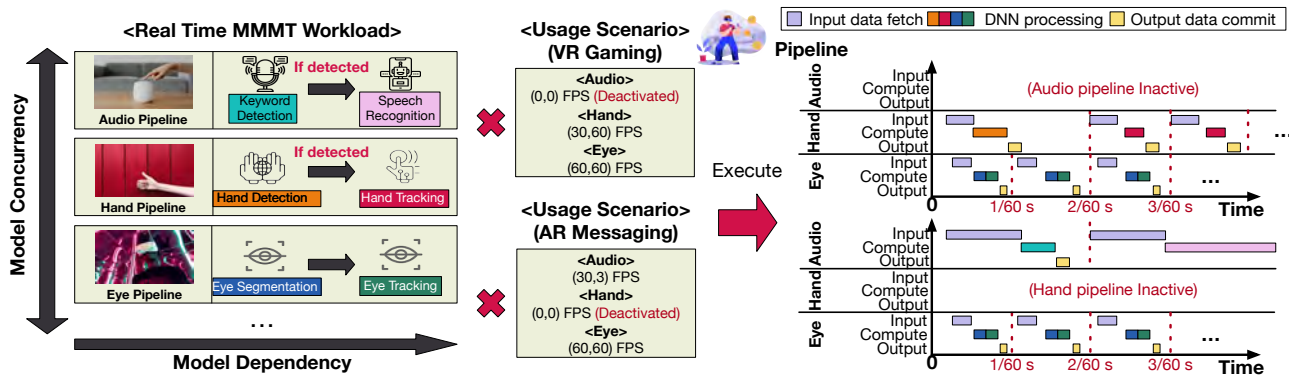


Figure 1. An example real-time multi-task multi-model (MTMM) ML workload and an example execution timeline.

Second, MTMM workloads exhibit *complex dependencies*. XR use cases display substantial data dependency (e.g., eye segmentation to tracking) and control dependency (e.g., hand detection to tracking) across models. These severe model-dependency limitations have ramifications for the underlying hardware and software scheduling space. In particular, the control flow dependencies make workload tasks dynamic, creating complexities for runtime scheduling.

Third, XR workloads have stringent user *quality of experience (QoE) requirements*. A key distinguishing factor of MTMM workloads from STSM ML workloads is the importance of understanding how to quantify the aggregated QoE metric across concurrent ML tasks at a system level. The resulting user QoE extends beyond the computational performance (latency or throughput) of a single model, which motivates the need for new metrics. Simple metrics like latency and/or FPS *do not* capture the complex interactions of all these models across diverse scenarios. For example, the latency of each inference cannot be the absolute metric, since improving latency beyond the deadline set by the target processing rate may not improve the overall processing rate (e.g., the processing rate may be bound by the sensor input stream rather than inference time). Therefore, we need a new scoring metric that can capture the aggregate performance of the MTMM workloads under different usage scenarios. The scoring metric must collectively consider all system aspects (model accuracy, achieved processing rate compared to the target processing rate, energy, etc.).

Collectively, not only do these three unique characteristics present system design challenges for XR, but they also make it challenging to benchmark and systematically characterize the performance of XR systems. Unfortunately, many of the characteristics and system-level concerns associated with MTMM workloads are not fully understood. This is largely due to the lack of public knowledge regarding the realistic characteristics of MTMM workloads derived from industry use cases. Consequently, the ML system design area for these workloads has yet to be explored. Furthermore, there is no benchmark suite of MTMM workloads that reflects industrial use cases. Many industry and academic benchmark

suites that exist today focus almost exclusively on STSM or MTMM without cascaded models (Reddi et al., 2020).

To address these deficiencies, we develop XRBENCH, a real-time multi-model ML benchmark with new metrics tailored for real-time MTMM workloads such as from the metaverse. XRBENCH includes proxy workloads based on real-world industrial use cases taken from production scenarios. These proxy workloads encapsulate the end-to-end properties of MTMM workloads at both the ML kernel and system levels, enabling the study of a vast design space.

XRBENCH includes scenario-based FPS requirements for ML use cases, which reflect the complex dependencies found in applications driving system-design research in a large organization invested in XR. It also presents representative QoE requirements for making system decisions. XRBENCH consists of many usage scenarios of a metaverse end-user device that combines various unit ML models with different target processing rates to reflect the dynamicity and real-time features of MTMM workloads. Furthermore, to enable comprehensive evaluations of ML systems using XRBENCH, we also propose and evaluate new scoring metrics that encompass four distinct requirements for the QoE of real-time MTMM applications: (1) the degree of deadline violations, (2) frame drop rate, (3) system energy consumption, and (4) model performance (e.g., accuracy).

In summary, we make the following contributions:

- We provide a taxonomy of MTMM-based workloads to articulate the unique features and challenges of real-time workloads for metaverse use cases.
- We present XRBENCH, an ML benchmark suite for real-time XR workloads. We provide open-source reference implementations for each of the models to enable widespread adoption and usage.
- We establish new scoring metrics for XRBENCH that capture key requirements for real-time MTMM applications and conduct quantitative evaluations.
- We make XRBENCH available as an open-source project: <https://github.com/XRBench>

Table 1. XRBENCH unit tasks and proxy unit models. Note that KD and SR are used for multiple task categories. Model performance requirements are 95% of model performance (or, 105% of error) reported in original papers, which opens the benchmark to various optimization techniques (e.g. mixed-precision), while ensuring reasonable prediction correctness. LT and GT refer to less than and greater than. For some models, we down-scale dataset resolution to adjust to the context of wearable/mobile devices, as we list in appendix A.

Category	Task	Model	Dataset	Model Perf. Requirement
Interaction	Hand Tracking (HT)	Hand Shape/Pose (Ge et al., 2019)	Stereo Hand Pose (Zhang et al., 2017)	AUC PCK, GT 0.948
	Eye Segmentation (ES)	RITNet (Chaudhary et al., 2019)	OpenEDS 2019 (Garbin et al., 2019)	mIoU, GT 90.54
	Gaze Estimation (GE)	Eyecod (You et al., 2022)	OpenEDS 2020 (Palmero et al., 2021)	Angular Error, LT 3.39
	Keyword Detection (KD)	Key-Res-15 (Tang & Lin, 2018)	Google Speech Cmd (Google, 2017)	Accuracy, GT 85.60
	Speech Recognition (SR)	Emformer (Shi et al., 2021)	LibriSpeech (Panayotov et al., 2015)	WER (others), LT 8.79
Context Understanding	Semantic Segmentation (SS)	HRViT (Gu et al., 2022)	Cityscape (Cordts et al., 2016)	mIoU, GT 77.54
	Object Detection (OD)	D2Go (Meta, 2022b)	COCO (Lin et al., 2014)	boxAP, GT 21.84
	Action Segmentation (AS)	TCN (Lea et al., 2017)	GTEA (Fathi et al., 2011)	Accuracy, GT 60.8
	Keyword Detection (KD)	Key-Res-15 (Tang & Lin, 2018)	Google Speech Cmd (Google, 2017)	Accuracy, GT 85.60
	Speech Recognition (SR)	Emformer (Shi et al., 2021)	LibriSpeech (Panayotov et al., 2015)	WER (others), LT 8.79
World Locking	Depth Estimation (DE)	MiDaS (Ranftl et al., 2020)	KITTI (Geiger et al., 2012)	$\delta > 1.25$, LT 22.9
	Depth Refinement (DR)	Sparse-to-Dense (Ma & Karaman, 2018)	KITTI (Geiger et al., 2012)	δ_1 , GT 85.5(100 samples)
	Plane Detection (PD)	PlaneRCNN (Liu et al., 2019)	KITTI (Geiger et al., 2012)	$AP^{0.6m}$, GT 0.37

2 MTMM WORKLOAD CHARACTERISTICS

To assist XR systems research on real-time MTMM workloads, we define a benchmark suite based on industrial meta-verse MTMM use cases. Before discussing the benchmark suite in Section 3, we first define the MTMM classification and the characteristics of a realistic MTMM workload, cascaded and concurrent MTMM.

2.1 Multi-model Machine Learning Workloads

Unlike STSM workloads, MTMM workloads include many models that lead to multiple model organization choices for constructing a workload instance. Based on the styles of those, we define three major classes:

- *Cascaded MTMM (cas-MTMM)*: Run multiple models back-to-back to enable one complex functionality (e.g., audio pipeline in Figure 1).
- *Concurrent MTMM (con-MTMM)*: Run multiple models independently at the same time to enable multiple unit functionalities (e.g., run Mask R-CNN (He et al., 2017) and PointNet (Qi et al., 2018) to perform 2D and 3D object detection during mapping and localization).
- *Cascaded and concurrent MTMM (cascon-MTMM)*: Hybrid of cas- and con-MTMM; connect multiple models back-to-back (cas-MTMM style) to implement a complex ML pipeline and deploy multiple models (con-MTMM style) for the other functionalities. (e.g., the VR gaming usage scenario in Figure 1).

Static vs. Dynamic: In addition to the model organization style, the model execution graph can be static or dynamic depending on the unit pipelines defined for a workload. For example, as shown in Figure 1, hand tracking can be deactivated if the hand detection model detects no hand.

In recent applications that encompass extended reality, we can observe dynamic and real-time cascon-MTMM style workloads (Kwon et al., 2021), which represent some of the most complicated ML inference workloads today. Although such dynamic and real-time cascon-MTMM style workloads are emerging, we lack a benchmark suite for dy-

namic cascon-MTMM workloads. Consequently, there has been no deep understanding of the features and challenges from dynamic cascon-MTMM, which we discuss next.

2.2 Dynamic Cascon-MTMM Features and Challenges

Cascaded and concurrent MTMM workloads are an emerging class of ML inference tasks. They have unique features and issues that do not exist in conventional ML workloads. We outline such aspects and analyze the issues of cascon-MTMM workloads for metaverse (XR) applications.

2.2.1 Scenario-driven Workloads

Metaverse workloads come from various different usage scenarios. A usage scenario refers to specific user experiences while utilizing a device or service. Gaming (e.g., VR gaming) and social (e.g., AR messaging) are example usage situations. Usage scenarios can be generated by combining several unit tasks, such as hand tracking or keyword detection. So, metaverse workloads must take the usage scenario into account to determine which unit tasks should be included, which is one of their distinctive elements compared to workloads in benchmarks such as MLPerf (Reddi et al., 2020) and ILLIXR (Huzafa et al., 2021).

2.2.2 Real-Time Requirements

Many existing ML-based applications often employ a single model inference to input (e.g., image or text). In contrast, metaverse devices are frequently required to continually execute inferences of a set of models in order to provide continuous user experiences (e.g., a user plays a VR game for 1 hour). As inference runs contribute to user experiences, it is only reasonable that a strong quality of user-driven experience (QoE) is required. In the context of multi-model inference, QoE can be represented by processing rate (i.e., inferences per second, such as FPS for models with frame-based inputs) or processing deadlines, hence introducing real-time processing requirements. Consequently, just as ML benchmarks must satisfy a certain level of accuracy for the quality of results (Reddi et al., 2020), XR benchmarks must also provide target processing rates.

2.2.3 Dynamic Cascading of Models

Metaverse applications commonly utilize numerous models. For example, hand-based interaction capabilities can be enabled by cascaded hand detection and hand tracking models. Such models are often cascaded (i.e., run sequentially in a back-to-back manner), and such cascaded models are characterized as a pipeline of models (or an ML pipeline).

Figure 1 presents three ML pipeline examples. Such pipelines need to be transformed into data dependencies across models, which need to be considered while scheduling computations (Kwon et al., 2021). MTMM ML pipelines may deactivate one or more downstream models based on the upstream model’s results. For instance, when no hand is detected, the hand tracking pipeline does not initiate the downstream hand tracking model. Such a dynamic aspect presents another problem for scheduling computation. In addition, it indicates that metaverse benchmarks must include different usage scenarios that reflect the dynamic nature of metaverse workloads.

2.2.4 Battery Life and Device Form Factor

The wearable form factor of metaverse devices makes thermal tolerances and battery life first-order priorities for user experience. For example, if the heat dissipation is excessively high, it may lead to skin discomfort or burns. Long battery life is critical since wearable devices are intended to be used continuously throughout the day, but the form factor places further constraints on battery size, even compared to other edge devices. For example, a recent metaverse device (goo, 2019) has an 800 mAh battery, which is a fifth of the size of the battery in a modern mobile device (e.g., 4000 mAh in Samsung Galaxy S20 (gal, 2019)). Energy consumption must be a primary optimization priority for metaverse end-user devices. All of the requirements (i.e., scenario-driven tasks, real-time requirements, dynamic cascading, battery life, and form factor) translate into energy constraints. So the benchmark needs to contain energy goals to ensure a device provides a good user experience.

3 XRBENCH

Real-time cascon-MTMM workloads for the metaverse are distinctive due to the discussed characteristics and obstacles. As a result, this domain necessitates a new method of defining benchmarks in comparison to traditional model-level benchmarks alone. In this section, we outline what we consider to be the most important characteristics of an MTMM benchmark. Then, we describe XRBENCH, the first benchmark of its kind for extended reality applications.

3.1 Benchmark Principles

To systematically guide the design of MTMM benchmarks, we focus on the key requirements for such a benchmark:

- *Usage Scenarios*: A set of real-world usage situations based on production use cases and a list of models to be run for each usage scenario must be defined.
- *Model Dependency*: As certain ML models are cascaded, model dependencies across the task must be specified to study resource allocation and scheduling effects.
- *Target Processing Rates*: Provide meaningful and applicable real-time requirements and processing rates for each model in each usage scenario to establish application behavior and system performance expectations.
- *Variants of a Usage Scenarios*: To reflect the dynamic nature of model execution and enable apples-to-apples comparisons, the benchmark must give numerous scenarios with distinct active time windows for each model.

Based on the requirements, we define XRBENCH. We first discuss unit models and usage scenarios in XRBENCH, then describe its evaluation infrastructure and scoring techniques. Later in Section 4, we show why these principles are important by conducting architectural analysis using XRBENCH.

3.2 Unit-level ML Models

Based on our experience in the metaverse (XR) domain, we define the first dynamic cascon-MTMM benchmark that reflects metaverse use cases. There are three main task categories in XRBENCH, listed in Table 1: real-time user interaction, understanding user context, and world locking (AR object rendering on the scene). These categories are based on real-world industrial use cases for the metaverse. For each unit task, we choose a representative reference model from the public domain.

When selecting models, we consider two aspects: (1) model performance (e.g., accuracy) and (2) efficiency (the number of FLOPs and parameters). Additionally, we list datasets and accuracy requirements for each unit task. More information for each unit task, including specific open-source model instances and dataset can be found in appendix A.

3.2.1 Interaction

Real-time user interaction tasks enable users to control metaverse devices using various input methods, including hand movements, eye gaze, and voice inputs. Therefore, we include corresponding ML model pipelines: hand pipeline (end-to-end model performing hand detection and tracking), eye pipeline (ES and GE), and voice pipeline (KD and SR).

3.2.2 Context Understanding

Context understanding tasks use multi- (e.g., VIO) or single-modal (e.g., audio) inputs to detect the context surrounding users so that a metaverse device can provide the appropriate user services. When a metaverse device detects that a user has entered a hiking trail, for example, it can provide the user with meteorological information. Context understanding models include scene understanding (SS, OD, and AS) and audio context understanding (KD and SR).

Table 2. Target processing rates (FPS). Eye and speech pipelines have data (D) or control (C) dependencies.

Usage Scenario	HT	Eye Pipeline		Speech Pipeline		SS	OS	AS	DE	DR	PD	Example Usage Scenario Description
		ES → GE (dep: D)	GE → ES (dep: C)	KD → SR (dep: C)	SR → KD (dep: D)							
Social Interaction A	30	60	60							30		AR messaging with AR object rendering
Social Interaction B		60	60					30				In-person interaction with AR glasses
Outdoor Activity A				3	3	10	30					Hiking with smart photo capture
Outdoor Activity B				3	3		30					Rest during hike
AR Assistant				3	3	10	10		30		30	Urban walk with informative AR objects
AR Gaming	45								30		30	Gaming with AR object
VR Gaming	45	60	60									Highly-interactive Immersive VR gaming

3.2.3 World Locking

A metaverse device must comprehend distances to real-world surfaces and occlusions in order to depict an augmented reality (AR) object on the display. These tasks are handled by models in the world-locking category, which includes a depth estimate model, a depth refinement model, and a plane detection model. The depth model is used to calculate the correct size of augmented reality (AR) objects, while the plane detection model identifies real-world surfaces that can be used to depict metaverse objects.

3.3 Usage Scenarios and Target Processing Rates

The models in Table 1 are selectively active with varying target processing rates depending on usage scenarios, as explained in Subsection 2.2.1. For example, the user experience of an AR game based on intensive hand interaction requires high hand-tracking speeds. The speech pipeline may be completely stopped if the game does not use speech input. During outdoor activities like hiking, however, an AR-enabled metaverse device may not require hand-tracking functionality but must be prepared for user speech input.

To reflect the different usage scenarios and target processing rate characteristics, we chose five realistic metaverse scenarios: (1) social interaction (AR messaging with AR object rendering), (2) outdoor activity (smart photo capture during hiking), (3) AR assistant (AR information display based on user contexts), (4) AR gaming, and (5) VR gaming.

Even within the same usage scenario, active models can differ because of the dynamic nature of cascon-MTMM workloads. For example, in an outdoor activity (hiking) usage scenario, when a user takes a break and tries to utilize an AR device (e.g., navigation and photo capturing), the hand tracking model will be engaged, unlike the previous hiking scenario. Considering such variability within usage scenarios, we suggest two versions (A and B) of social interaction and outdoor activity scenarios. Table 2 describes the usage scenario variants. In addition, we specify a target processing rate for each model with three levels: High (60Hz or 45Hz), Medium (30Hz), and Low (10Hz). SR has a processing rates of 3Hz, which models the 320ms left context size utilized in its original paper (Shi et al., 2021). We assign target processing speeds depending on the usage scenarios based on practical metaverse use cases. We suppose that a metaverse device identifies the active models and their

Table 3. Three main input sources to a metaverse device. We align all the input streaming rates to be 60 FPS for multi-modal models (e.g., DR in Table 1). We also model jitters for each data frame.

Input Source	Input Type	Streaming Rate	Jitter
Camera	Images	60 FPS	$\pm 0.05 ms$
Lidar	Sparse Depth Points	60 FPS	$\pm 0.05 ms$
Microphone	Audio	3 FPS	$\pm 0.1 ms$

processing rates (i.e., usage scenario) based on the specific application launched by a user.

3.4 Input Sources and Load Generation

Metaverse devices utilize multiple sensors with varying modalities. To model the sensors, we use the settings listed in Table 3 for the unit models in Table 1. The camera is the input source of images used by computer vision models. The lidar sensor provides a sparse depth map to the depth refinement model using RGBd data. The microphone receives audio inputs for speech models (KD and SR).

The arrival time of the input data in an actual system can vary slightly from the projected time based on the streaming rate, depending on multiple circumstances (e.g., system bus congestion). In numerous research analyses, jitter is frequently disregarded. However, in genuine production usage circumstances, jitter might result in sporadic frame dropouts, which degrades QoE. To represent such aspects, we apply a jitter to each data frame, as shown in Table 3, and we alter the injection time of inference requests accordingly.

3.5 Benchmark Harness

A harness orchestrates the execution of the models, respecting the dependencies. We illustrate the structure of the benchmark harness in Figure 2. The harness takes workload and system information as input, and generates reports that contain not only the scores (overall score and its breakdowns; to be discussed in Subsection 3.7) but also detailed performance statistics such as the amount of delay over deadline, frame drop, execution timeline, and so on. We include this detailed information in the reports to help users use XRBENCH to guide their system designs.

The harness consists of a runtime, logger, and scoring module. The runtime contains a load generator which intermittently generates jittered inference requests. The inference dispatcher/scheduler is the core component of the runtime, which (1) selects the next inference requests to be dispatched when a hardware entity (e.g., accelerator) becomes available, (2) tracks the model and frame dependencies, and (3)

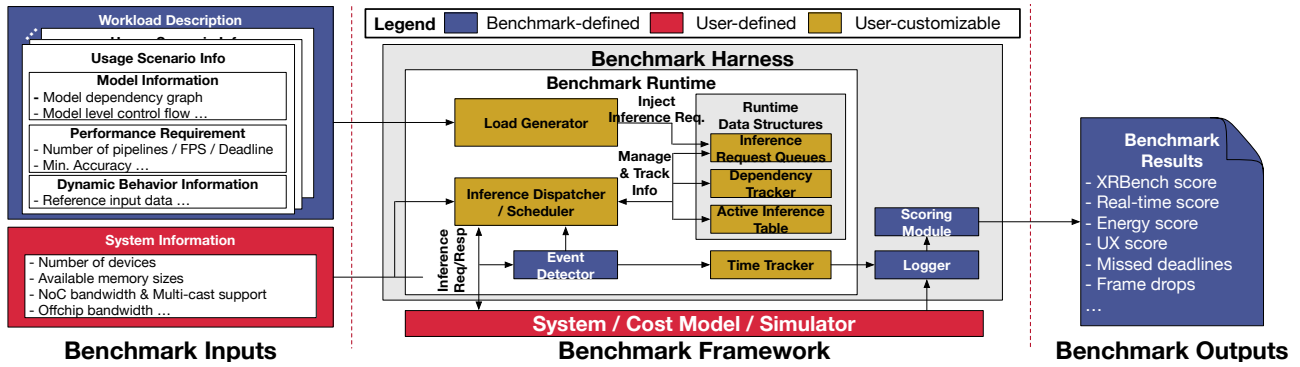


Figure 2. An overview of the benchmark harness, XRBENCH.

dispatches inferences to the machine learning system to be evaluated (which may be a real physical system, analytical cost model, or simulator). The runtime components include an event detector, score tracker, and various data structures (request queue, active inference table, dependency table, etc.) that assist the dispatcher and scheduler.

XRBENCH requires users to finish a certain number of runs, which equals to the target processing rate within a set duration (default: one second), to ensure the real-time requirement is satisfied. XRBENCH provides a simple latency-greedy (for cost model or simulator-based runs) or a round-robin style scheduler (for real systems) for models within each usage scenario. Users can replace the scheduler and other components highlighted in yellow boxes in Figure 2 to model their runtime or system software’s behavior. Much like in a traditional ML system or ML benchmark (Reddi et al., 2020), optimizing the software stack is crucial to the hardware’s success, and XRBENCH encourages such optimizations.

3.6 Deep Dive Example

To clarify the roles of each piece in XRBENCH, Figure 3 provides an example execution timeline for the “Social Interaction A” usage scenario in Table 2. The execution graph on the left shows the active models, their processing rates, and the model dependencies. The right side corresponds to a sample execution timeline for the scenario. The top-most section represents the input streaming from relevant input sources listed in Table 3. Each input source can have different initial delays and jitter.

A compute engine (such as an accelerator) can only execute model inferences if it has access to the input data. In this example, we model a simple scheduler assuming that inferences can only begin if the input data is ready. Consequently, Eye Segmentation (ES) and Gaze Estimation (GE) for frame 0 begin once the input data retrieval for image frame 0 concludes. Additionally, GE runs after ES to satisfy their dependency. The multi-modal Depth Refinement (DR) model executes after image and depth point inputs are received. As DR’s target processing rate is 30 FPS while depth point input streams at 60 FPS, only every other input

is used. As Hand Tracking (HT) also operates at 30 FPS, it skips every other image frame. The DR output is used for display-targeted AR object rendering. Therefore, DR results must be delivered by a certain time, which is a 30 FPS deadline (e.g., $2/60$ s for frame 0) in this example.

In the execution timeline, the usage scenario is effectively supported and desired processing rates are attained. However, ET and HT results are delivered past their desired deadlines ($1/60 \times frame$ and $1/30 \times frame$ seconds, respectively). This suggests that HT and ET latency must be reduced further in this example. Not always should latency itself be used as an optimization target, though, as latency reduction beyond deadlines may not improve the user experience. Even zero-latency inferences cannot increase the effective processing rate of a task beyond the input data streaming rates because future data cannot be processed without them. This raises the question: How should we quantify the performance of a real-time MTMM system by taking into account the actual quality of the results for users? In the next subsection, we discuss new metrics that encompass the requirements and characteristics of XR tasks.

3.7 Scoring Metrics

In Subsection 3.6, we showed that evaluating a system for real-time MTMM workloads is not trivial using the example in Figure 3. For example, lower inference latency does not always improve user experiences if the processing rate of each task is bound by the input data streaming rate. We need to capture such aspects when we evaluate a system for real-time XR workloads. Therefore, we define a new scoring metric, XRBENCH SCORE, considering all the aspects we discussed and propose it to be used as the overall performance metric in XRBENCH.

Based on the unique features of XR workload, we list the following score requirements for the benchmark:

- [Real-time] The score should include a penalty if the latency exceeds the usage scenario’s required performance constraints (i.e., missed deadlines).
- [Low-energy] The score should prioritize low-power designs as metaverse devices are energy-constrained.

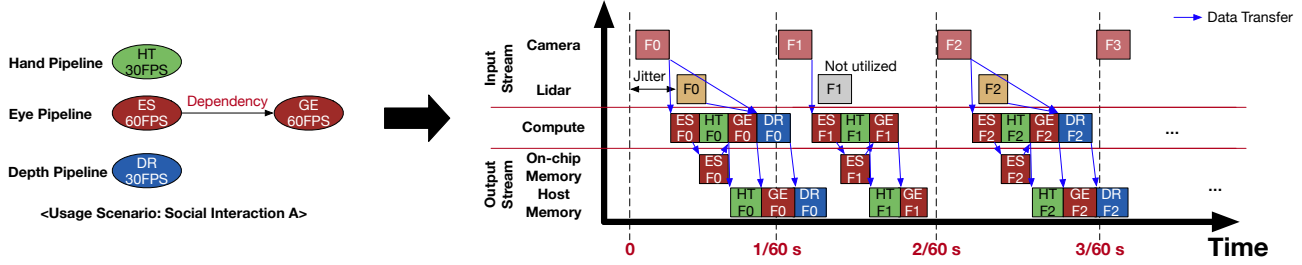


Figure 3. An example execution timeline based on the Social Interaction A usage scenario in Table 2, FN refers to the frame N.

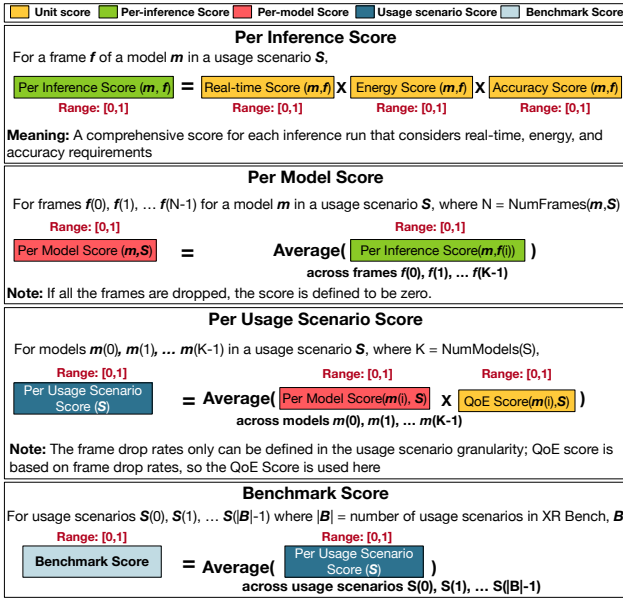


Figure 4. A high-level overview of how we define benchmark scores at inference run, model, and usage scenario granularity using unit scores (real-time, energy, accuracy, and QoE scores).

- [Model quality] The score should capture the output quality delivered to a user from running all the models.
- [QoE requirement] The score should include a penalty if the FPS drops below the target FPS to maintain QoE.

We define four unit scores: real-time (RT), energy, accuracy, and QoE scores. Each score is constrained to be in the [0, 1] range for easy breakdown analysis. We multiply unit scores to reflect all of their aspects while keeping the results in the [0,1] range. We utilize averages to summarize scores for multiple inference runs on different frames for a model, multiple models within a usage scenario, etc. We focus on high-level intuitions with detailed formal definitions presented in Table 4, Box 1, and Box 2.

To model real-time requirements, we consider the following observations: (1) too much optimization on inference latency beyond the deadline does not lead to higher processing rates. (2) reduced latency can still be helpful for scheduling other models. (3) violated deadlines gradually disrupt the user experience (e.g., Achieving 59 FPS for an eye-tracking model targeting 60 FPS won't significantly

Table 4. Symbols used in the formulation. (Only listing those not defined in Box 1 and Box 2.

Symbol	Definition
M_{ID}	Model descriptor (model name)
$inSrcID$	Input source descriptor (e.g., sensor)
DS_{ID}	Dataset descriptor
QM_{ID}	Model quality metric descriptor (e.g., accuracy)
QM_{targ}	Target value of a model quality metric
QM_{Type}	The type of QM. Either Higher- or lower-is-better (HiB/LiB)
J_t	Max absolute jitter in ms ($J_t \geq 0$)
Dep_{μ}	A set of models on which model μ depend
L_{init}	Initialization latency (ms) of an input stream
L_{inf}	Latency for an inference run
$T_{start}(h, \mu)$	Start time of an inference of μ on hardware h
$T_{end}(h, \mu)$	Completion time of an inference of μ on hardware h
HiB/LiB	Higher-is-better and lower-is-better metrics

affect the user experience). Based on these observations, we search for a function that (1) gradually rewards/penalizes for reduced/increased latency near a deadline and (2) outputs 0 and 1 if the latency is well beyond (e.g., 0.5ms for a deadline of 10ms) or within the deadline, respectively. We find such a function by modifying the sigmoid function, which is widely used in ML models.

For energy, a lower-is-better metric, a naive way to compute energy score is computing the inverse of the energy consumption (example unit: 1/mJ). However, the range of the naive metric is unbounded, which makes it hard for component-wise analysis when it is combined with other scores bound in [0,1] ranges. Therefore, to bound the energy score within the [0,1] range as well, we utilize a large energy limit E_{max} to define the top-end of the score.

For accuracy score, we quantify how much the output correctness differs from the desired level using model-specific performance metrics (e.g., accuracy for classification, mIoU for segmentation, PCK AUC for hand tracking, etc.). Although there are many different metrics other than accuracy, we use the term, accuracy score, for simplicity.

Finally, we construct the XRBENCH SCORE in a hierarchical manner. Figure 4 illustrates how we combine scores along stages (unit, per-inference, per-model, per-usage scenario) to finally generate the overall XRBENCH SCORE. We first compute the per-inference score by multiplying real-time, energy, and accuracy scores. The QoE score is not applied here as the frame drop rate only can be defined at the usage scenario level since the FPS requirements change depending on the usage scenario. Using the per-inference

System/Benchmark Parameters

$$\begin{aligned}
M_{ID}, inSrcID, DS_{ID}, QM_{ID} &\in str \\
FPS_{sensor}, FPS_{model}, InFrame_{ID} &\in \mathbb{N} \\
L_{init}, L_{inf}, Jt, QM_{targ}, T_{req}, \epsilon &\in \mathbb{R} \\
QM_{Type} &= HiB \mid LiB
\end{aligned}$$

Input Data Stream (St_{input})

$$St_{input} = \{\sigma \mid \sigma = (inSrcID, FPS_{sensor}, L_{init}, Jt)\}$$

Model Quality Goal (Q)

$$Q = (QM_{ID}, QM_{Targ}, QM_{Type})$$

Unit Models (M)

$$M = \{\mu \mid \mu \in (M_{ID}, DS_{ID}, \sigma, Q) \wedge \sigma \in St_{input}\}$$

Usage Scenario (θ)

$$\theta = \{(\mu, Dep_{\mu}, FPS_{model}) \mid \mu \in M \wedge Dep_{\mu} \subset M\}$$

Benchmark Suite (Ω)

$$\Omega = \{\theta_1, \theta_2, \dots, \theta_{NumScn}\}$$

Inference Request (IR)

$$IR = (\mu, InFrame_{ID})$$

Inference Request Time ($T_{req}(IR)$)

$$\begin{aligned}
T_{req}(IR) &= L_{init}(inSrcID) + \frac{InFrame_{ID}}{FPS_{sensor}(inSrcID)} \\
&+ 2Jt (Dist(rand(inSrcID \times InFrame_{ID})) - 0.5) \\
&\text{where } Dist(x) \in [0, 1] \wedge x \in \mathbb{R}
\end{aligned}$$

Inference Deadline ($T_{dl}(IR)$)

$$T_{dl}(IR) = L_{init}(inSrcID) + \frac{InFrame_{ID} + 1}{SR(inSrcID)}$$

Inference Slack ($T_{sl}(IR)$)

$$T_{sl}(IR) = T_{dl}(IR) - T_{req}(IR)$$

Box 1. Base Definitions

score, we construct the per-model score by computing the average across all processed frames. We do not include dropped frames since they will be considered in the QoE score. To compute the per-usage scenario score, we compute the average of the product of per-model score and QoE score across all the models within a usage scenario.

Based on our approach discussed in this section, we formalize our score metrics in Table 4, Box 1, Box 2. We also provide more details in appendix B.

XR BENCH reveals all individual scores to users to facilitate Pareto frontier analysis, in addition to XR BENCH SCORE. In some cases, the industry may not wish to share the detailed performance breakdown of their system. Therefore, reporting breakdown scores is optional for XR BENCH, while the overall XR BENCH SCORE is mandatory. The released benchmark harness contains implementations of all scoring metrics.

3.8 Limitations

XR BENCH focuses on ML components of XR workloads and does not model pre- and post-processing of inputs and outputs of ML pipelines. Such an approach is motivated by the significance of the ML processing time in XR systems.

Unit Score: Realtime Score ($RtScore(IR)$)

$$RtScore(IR) = \frac{1}{1 + e^{k(L_{Inf}(IR) - T_{sl}(IR))}}$$

Unit Score: Energy Score ($EnScore(IR)$)

$$EnScore(IR) = \frac{En_{max} - En(IR)}{En_{max}}$$

Unit Score: Accuracy Score ($AccScore(IR)$)

$$AccScore(IR) = \max(1, rawAccScore(IR))$$

$$rawAccScore(IR) = \begin{cases} \frac{QM_{measured}}{QM_{targ}}, & \text{if } QM_{Type} = HiB \\ \frac{QM_{measured}}{QM_{measured} + \epsilon}, & \text{otherwise} \end{cases}$$

where $\epsilon > 0 \wedge \epsilon \ll 1 \wedge \epsilon \in \mathbb{R}$

Unit Score: QoE Score ($QoEScore(\mu)$)

$$QoEScore(\mu) = \frac{NumFrm_{exec}(\mu)}{NumFrm(\mu)}$$

Aggregated Score: Inference-wise Score ($Score_{inf}(IR)$)

$$\begin{aligned}
Score_{inf}(IR) &= RtScore(IR) \times EnScore(IR) \\
&\times AccScore(IR)
\end{aligned}$$

Aggregated Score: Usage Scenario Score ($Score_{scn}(\theta)$)

$$Score_{scn}(\theta) = \sum_{j=1}^{NumFrm(\mu)} \frac{Score_{inf}(IR) \times QoEScore(\mu)}{NumFrm(\mu) \times |\theta|}$$

Aggregated Score: XRBENCH Score ($Score_{bench}$)

$$Score_{bench} = \frac{\sum_{\theta \in \Omega} Score_{scn}(\theta)}{|\Omega|}$$

Box 2. Score metrics

4 EVALUATION

In this section, we focus on three key questions to ascertain the value of XR BENCH: (1) why the comprehensive overall score is necessary for the proper evaluation of XR tasks, (2) why it is important to study the different usage scenarios that are included in XR BENCH, and (3) what are the hardware implications of the MTMM characteristics found in XR.

4.1 Methodology

Metaverse applications run on wearable devices and the compute requirement for the workloads is heavy (tens of FPS requirements for multiple models). Therefore, considering the capabilities of state-of-the-art mobile SoCs (e.g., 26 TOPS on Qualcomm Snap Dragon 888 (Qualcomm, 2022)), we model wearable devices with DNN inference accelerators that employ 4K and 8K PEs with 256 GB/s on-chip bandwidth and 8MiB of on-chip shared memory running at 1 GHz clock, similar to Herald (Kwon et al., 2021).

Simulated HW Accelerators. Table 5 shows various accelerator instances we evaluate in three accelerator styles: **FDA** (fixed-dataflow accelerator), **Scaled-out multi-FDA** (two accelerator instances with the same dataflow style motivated by (Baek et al., 2020)), and **HDA** (heterogeneous dataflow accelerator) (Kwon et al., 2021). Depending on the style, we partition the 4K and 8K PEs into 2 or 4 accelerator instances. The WS (weight stationary) dataflow is

Table 5. Accelerator styles. Partitioning indicate the PEs to be deployed for each accelerator instance for SFDA and HDA.

Acc. ID	Acc. Style	Dataflow
A	FDA	WS
B		OS
C		RS
D	SFDA ¹	WS + WS (1:1 partitioning)
E		OS + OS (1:1 partitioning)
F		RS + RS (1:1 partitioning)
G		WS + WS + WS + WS (1:1:1:1 partitioning)
H		OS + OS + OS + OS (1:1:1:1 partitioning)
I	RS + RS + RS + RS (1:1:1:1 partitioning)	
J	HDA	WS + OS (1:1 partitioning)
K		WS + OS (3:1 partitioning)
L		WS + OS (1:3 partitioning)
M		WS + OS + WS + OS (1:1:1:1 partitioning)

inspired by NVDLA (NVIDIA, 2017) that parallelizes the output and input channels with input columns. OS (output-stationary) is a hand-optimized dataflow that parallelizes output rows and columns with a 16-way adder tree reducing input channel-wise partial sums. The RS (row stationary) dataflow is inspired by Eyeriss (Chen et al., 2016) that parallelizes output channels, output rows, and kernel rows. Note that each accelerator in Table 5 refers to an instance of hardware accelerator that can run XRBENCH.

Simulation Methodology. We implement the framework illustrated in Figure 2 and plug in MAESTRO (Kwon et al., 2019) as the analytical cost model to perform the different case studies. All the models are the same across the hardware platforms (8bit-quantized without other optimizations) and satisfy the accuracy goals (i.e., accuracy score = 1).

Modeling Dynamic Cascading. To model the dynamic cascading between keyword detection and speech recognition, we apply pre-defined probabilities of user keyword utterances to corresponding usage scenarios (Outdoor A, Outdoor B, and AR Assistant). For outdoor activity scenarios, we apply 0.2 as the interaction is expected to be in a low frequency for the scenarios. For AR assistant, we apply 0.5 as the speech is the standard interaction method for the use case. For eye segmentation and gaze estimation pipeline, we first apply the probability of 1.0 to model pure data dependency and sweep the probability for a separate deep dive (Figure 7).

4.2 Why the XRBENCH SCORE is a Necessary Metric

The intent of this section is to show that the overall scoring metric we present (Section 3.7) is necessary for systematically evaluating XR systems. We present our evaluation results in Figure 5, which shows score break-downs for each accelerator style running each usage scenario.

4.2.1 Overall Score Enables Comprehensive Analysis

The real-time score quantifies the degree of deadline violation. Higher-is-better for the real-time score; however, a high real-time score itself does not guarantee ideal system performance. For example, accelerator A with 8K PEs run-

ning the Outdoor Activity B (Figure 5, (d)) has a real-time score of 1.0, which indicates that most of the deadlines are met within a small margin. However, accelerator A misses 10.0% of the frames (not shown) and has high energy consumption, 34.1% greater than the most energy-efficient design (accelerator C). Our scoring metric incorporates all aspects, including QoE score for frame drops and energy score for energy consumption, and it reports an overall score of 0.49, which is 42.9% less than the best accelerator (I).

As another example, for the AR Gaming scenario (Figure 5, (g)) on a 4K accelerator system, accelerator G achieved the greatest QoE score of 0.91 and a strong energy score of 0.76. However, its real-time score is zero due to heavily missed deadlines. That is, while the frame rate is overall close to the target as captured in the QoE score, a user will experience heavy output lag, which degrades the real-time experience. The real-time score captured this and led the overall score for this accelerator to be zero.

4.2.2 Hardware Utilization is the Wrong Metric

Hardware utilization is often used as a key metric for accelerator workloads since it can be directly translated to accelerator performance by multiplying utilization by the peak performance of the accelerator. However, we do not consider hardware utilization to be the right metric for real-time MTMM applications, and as such we do not include it in the overall scoring metric (Section 3.7).

Utilization does not consider frame drops or periodic workload injection. For example, Figure 6 shows the execution timelines for the 4K and 8K PE versions of accelerator J. The 8K PE timeline (Figure 6, (b)) has more gaps than the 4K PE timeline, which means the overall accelerator utilization is lower than that of the 4K PE accelerator (Figure 6, (a)), making it seem as though the 4K PE accelerator is a better choice. However, the 4K-PE accelerator drops 47.1% of the frames and completely fails to run the PD model, whereas the 8K-PE accelerator drops only 2.3% of frames.

Unlike the utilization alone, our score metrics properly capture the real-time requirement and QoE aspects. The frame drop rates of the 4K and 8K PE accelerators are captured in the QoE scores of 0.53 and 0.97, respectively. In addition, the large amount of deadline violations for the PD model in the 4K PE accelerator results in a real-time score of 0. Combining those unit scores into the overall score (XRBENCH SCORE), we observe the scores of 0 and 0.51, which provides a better intuition to the comprehensive performance of an XR system considering all the considerations, including real-time requirements and QoE.

4.3 Why It is Important to Dive into Usage Scenarios

Even though all usage scenarios in XRBENCH reflect the metaverse domain, the individual workload characteristics

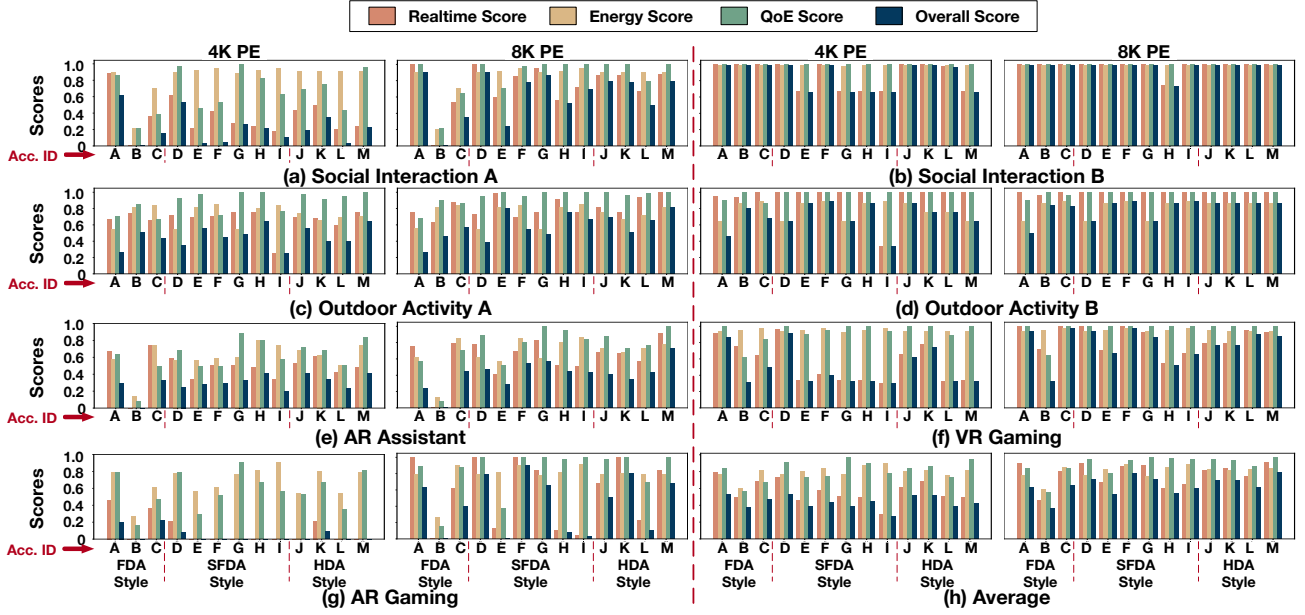


Figure 5. The scores computed for each style of an accelerator system with 4K and 8K PEs. (a-g) the score break-downs for each usage scenario. (h) the average across scenarios. Overall score refers to XRBENCH SCORE.

Table 6. List of existing benchmarks related to ML and XR workloads, with comparison of workload characteristics and score metrics. Δ means the property is partially supported by the benchmark.

Benchmark		Workload Characteristics				Score Metrics				
		Cascon-MMMT	Dynamic Workload	Real-time Scenarios	Focus on ML	Device Scope	Latency	Energy	Accuracy	QoE
ML	MLPerf Inference ^a			✓	✓	server	✓	✓	✓	
	MLPerf Tiny ^b			✓	✓	edge	✓	✓	✓	
	MLPerf Mobile ^c				✓	mobile	✓		✓	
	DeepBench ^d				✓	server/edge	✓			
	AI Benchmark ^e				✓	mobile	✓			
	EEMBC MLMark ^f				✓	edge	✓		✓	
	AlBench ^g	✓	Δ	✓	✓	server	✓	✓	✓	✓
XR	AloTBench ^h			✓	✓	mobile/edge	✓	✓	✓	
	ILLIXR ⁱ	✓	Δ	✓	✓	edge	✓	✓	Δ	✓
ML + XR	XRBENCH	✓	✓	✓	✓	edge	✓	✓	✓	✓

References: ^a (Reddi et al., 2020), ^b (Banbury et al., 2021), ^c (Reddi et al., 2022), ^d (Dee, 2016), ^e (Ignatov et al., 2018), ^f (EEM, 2020), ^g (Gao et al., 2019), ^h (Luo et al., 2018), ⁱ (Huzaifa et al., 2021), ^j (VRM, 2020)




Figure 6. Execution timeline of AR gaming scenario on 4k and 8k PE versions of WS and OS HDA accelerator (accelerator J).

are diverse and tend to vary during execution, resulting in different system performance. Each usage scenario prefers different accelerator types, as shown in Figure 5. For example, in the 4K PE config, the Social Interaction A scenario (Figure 5, (a)) prefers the FDA style accelerator with WS dataflow (accelerator A). However, Outdoor Activity A (Fig-




Figure 7. Evaluated scores on accelerators B and J with 4K PEs running VR gaming scenario. We vary the probability of triggering GE after ES, modeling the dynamic cascading.

ure 5, (c)) prefers the SFDA style with four sub-accelerators with the OS dataflow (accelerator H).

Moreover, dynamically cascaded models (Section 2.2) require a deep dive into corresponding usage scenarios. To understand the impact of dynamically cascaded models, we

vary the probability of triggering the GE model after the ES (assuming that GE is triggered only if ES results have sufficiently large segmented eyes). We run 200 experiments and plot the average data to capture the overall trend, focusing on low- and high-score cases (accelerators B and J) in Figure 7.

Overall, both designs maintain their overall scores while we observe a slight decline (0.03) in the overall score on the high-score design (accelerator J) as moving from 25% to 100% cascading probability. As the cascading probability increases from 25% to 100%, the real-time score of accelerator B increased by 0.23 points while the QoE score decreased by 0.06 points. This indicates that the low-score design (accelerator B) can drop some frames to maintain overall user experience quantified by XRBENCH SCORE under high cascading probability. Such results motivate further investigation of optimization opportunities in the scheduler and runtime for XR ML systems.

4.4 What the Implications to Future XR Systems Are

We list three observations we make from the evaluation.

Observation 1) XR systems need to be co-designed with usage scenarios. Evaluation results show that every usage scenario prefers different XR systems. For example, comparing accelerator styles in the 4K PE setting, we find the accelerator styles with the highest score are all different for each workload. For example, accelerator A (FDA style, WS dataflow, single-accelerator system) is the best style for the social interaction A scenario (Figure 5, (a)). However, accelerator F (SFDA style, OS dataflow, four-accelerator system) performed the best for the outdoor activity B scenario (Figure 5, (d)). Such results suggest that the XR systems require careful co-design with the usage scenarios.

Observation 2) Optimal accelerator styles depend on the chip size. The style H (SFDA style, RS dataflow, four-accelerator system) performs the best for the AR assistant scenario (Figure 5, (e)) with 4K PEs. However, when the total number of PEs changes to 8K, the style M (HDA style, WS and OS dataflows, four-accelerator system) performs the best. Those results imply that the design space for XR applications is complex with distinctive features of real-time MTMM workloads, which motivates follow-up studies using XRBENCH.

Observation 3) Multi-accelerator systems are friendly to XR workloads. We also find the preference of the number of models in MTMM models to the multi-accelerator system (e.g., SFDA and HDA). AR assistant (Figure 5, (e)) and VR gaming (Figure 5, (f)) scenarios include the most (6) and least (3) number of models, respectively. For AR assistant, we observe the multi-accelerator style (SFDA and HDA) outperforms the single accelerator style. For VR gaming scenario, in contrast, the FDA style (accelerator A)

outperforms most of the other accelerators. In particular, when the sub-accelerator size is sufficiently large (8K PE), a quad-accelerator system (HDA accelerator M) performs the best on the many-model scenario (AR assistant), but the same system underperforms on the fewer-model scenario (VR gaming). Such data show the efficacy of parallel model execution using sub-accelerators, which motivates to explore scale-out designs for many-model MTMM workloads like the AR assistant.

5 RELATED WORK

Based on the characteristics we describe in Section 3, we present the limitations of existing ML and XR benchmarks in Table 6. XRBENCH is unique in that it is the only workload suite that captures complex workload dependencies, is ML-focused, presents several real-world usage scenarios that are distilled from industry practice and uniquely establishes a robust scoring metric. Due to space limitations, we defer detailed discussions of the benchmarks to appendix C. In summary, XRBENCH is the first suite to include several ML workloads tailored for XR applications.

6 CONCLUSION

Metaverse use cases necessitate complex ML benchmark workloads that are essential for fair and useful analyses of existing and future system performance, but such workloads exceed the capabilities of existing benchmark suites. The XR benchmark we present, which is based on industry experience, captures the diverse and complex characteristics of these emerging ML-based MTMM workloads. We believe XRBENCH will foster new ML systems research focused on XR.

ACKNOWLEDGEMENTS

This work was enabled in part by support from Robert Shearer at Meta. We appreciate Rob’s ongoing assistance and counsel in helping us create methodical approaches to assess XR SoC designs. Part of the funding for authors from Harvard University came from the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) and the Semiconductor Research Corporation (SRC). The Georgia Tech authors were funded in part by SRC. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, SRC, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

- Deepbench. <https://github.com/baidu-research/DeepBench>, 2016.
- Ed-tn. https://github.com/colincls1/TemporalConvolutionalNetworks/blob/master/code/TCN_main.py, 2016.
- Aiotbench, benchcouncil. <https://www.benchcouncil.org/aibench/aiotbench/index.html>, 2018.
- Ritnet. <https://github.com/AayushKrChaudhary/RITnet>, 2019.
- Galaxy s20 specifications. https://www.samsungmobilepress.com/mediareources/galaxy_s20/techspecs, 2019.
- Glass enterprise edition 2 specifications. <https://www.google.com/glass/tech-specs/>, 2019.
- Eembc mlmark v2.0. <https://www.eembc.org/mlmark/>, 2020.
- Vrmark. <https://benchmarks.ul.com/vrmark>, 2020.
- midas_v21_small. https://github.com/AlexeyAB/MiDaS/releases/download/midas_dpt/midas_v21_small-70d6b9c8.pt, 2020.
- Hrvit-b1. <https://github.com/facebookresearch/HRViT/blob/main/models/hrvit.py#L1125-L1155>, 2022.
- Baek, E., Kwon, D., and Kim, J. A multi-neural network acceleration architecture. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 940–953. IEEE, 2020.
- Banbury, C., Reddi, V. J., Torelli, P., Holleman, J., Jeffries, N., Kiraly, C., Montino, P., Kanter, D., Ahmed, S., Pau, D., et al. Mlperf tiny benchmark. *arXiv preprint arXiv:2106.07597*, 2021.
- Barham, P., Chowdhery, A., Dean, J., Ghemawat, S., Hand, S., Hurt, D., Isard, M., Lim, H., Pang, R., Roy, S., Saeta, B., Schuh, P., Sepassi, R., Shafey, E. L., Thekkath, A. C., and Wu, Y. Pathways: Asynchronous distributed dataflow for ml. *Proceedings of Machine Learning and Systems*, 4:430–449, 2022.
- Chaudhary, A. K., Kothari, R., Acharya, M., Dangi, S., Nair, N., Bailey, R., Kanan, C., Diaz, G., and Pelz, J. B. Ritnet: Real-time semantic segmentation of the eye for gaze tracking. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pp. 3698–3702. IEEE, 2019.
- Chen, Y.-H., Emer, J., and Sze, V. Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks. In *International Symposium on Computer Architecture (ISCA)*, 2016.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3213–3223, 2016.
- Farrell, S., Emani, M., Balma, J., Drescher, L., Drozd, A., Fink, A., Fox, G., Kanter, D., Kurth, T., Mattson, P., et al. Mlperf™ hpc: A holistic benchmark suite for scientific machine learning on hpc systems. In *2021 IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC)*, pp. 33–45. IEEE, 2021.
- Fathi, A., Ren, X., and Rehg, J. M. Learning to recognize objects in egocentric activities. In *CVPR 2011*, pp. 3281–3288. IEEE, 2011.
- Gao, W., Tang, F., Wang, L., Zhan, J., Lan, C., Luo, C., Huang, Y., Zheng, C., Dai, J., Cao, Z., Zheng, D., Tang, H., Zhan, K., Wang, B., Kong, D., Wu, T., Yu, M., Tan, C., Li, H., Tian, X., Li, Y., Shao, J., Wang, Z., Wang, X., and Ye, H. Aibench: An industry standard internet service ai benchmark suite. *ArXiv*, abs/1908.08998, 2019.
- Garbin, S. J., Shen, Y., Schuetz, I., Cavin, R., Hughes, G., and Talathi, S. S. Openeds: Open eye dataset. *arXiv preprint arXiv:1905.03702*, 2019.
- Ge, L., Ren, Z., Li, Y., Xue, Z., Wang, Y., Cai, J., and Yuan, J. 3d hand shape and pose estimation from a single rgb image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Geiger, A., Lenz, P., and Urtasun, R. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 3354–3361. IEEE, 2012.
- Google. Google speech commands. <https://ai.googleblog.com/2017/08/launching-speech-commands-dataset.html>, 2017.
- Gu, J., Kwon, H., Wang, D., Ye, W., Li, M., Chen, Y.-H., Lai, L., Chandra, V., and Pan, D. Z. Multi-scale high-resolution vision transformer for semantic segmentation.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12094–12103, 2022.
- He, K., Gkioxari, G., Dollár, P., and Girshick, R. Mask rcnn. In *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.
- Huzaifa, M., Desai, R., Grayson, S., Jiang, X., Jing, Y., Lee, J., Lu, F., Pang, Y., Ravichandran, J., Sinclair, F., Tian, B., Yuan, H., Zhang, J., and Adve, V. S. Illixr: Enabling end-to-end extended reality research. 2021.
- Ignatov, A., Timofte, R., Chou, W., Wang, K., Wu, M., Hartley, T., and Van Gool, L. Ai benchmark: Running deep neural networks on android smartphones. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 0–0, 2018.
- Kwon, H., Chatarasi, P., Pellauer, M., Parashar, A., Sarkar, V., and Krishna, T. Understanding reuse, performance, and hardware cost of dnn dataflow: A data-centric approach. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pp. 754–768, 2019.
- Kwon, H., Lai, L., Pellauer, M., Krishna, T., Chen, Y.-H., and Chandra, V. Heterogeneous dataflow accelerators for multi-dnn workloads. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 71–83. IEEE, 2021.
- Lea, C., Flynn, M. D., Vidal, R., Reiter, A., and Hager, G. D. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 156–165, 2017.
- 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*, pp. 740–755. Springer, 2014.
- Liu, C., Kim, K., Gu, J., Furukawa, Y., and Kautz, J. Planercnn: 3d plane detection and reconstruction from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4450–4459, 2019.
- Luo, C., Zhang, F., Huang, C., Xiong, X., Chen, J., Wang, L., Gao, W., Ye, H., Wu, T., Zhou, R., and Zhan, J. Aiot bench: Towards comprehensive benchmarking mobile and embedded device intelligence. In *Bench*, 2018.
- Ma, F. and Karaman, S. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 4796–4803. IEEE, 2018.
- Meta. Fbnet-c. <https://github.com/facebookresearch/mobile-vision>, 2019.
- Meta. Faster-rcnn-fbnetv3a. https://github.com/facebookresearch/d2go/blob/main/configs/faster_rcnn_fbnetv3a_C4.yaml, 2022a.
- Meta. D2go. <https://github.com/facebookresearch/d2go>, 2022b.
- Meta. What is the metaverse? <https://about.facebook.com/what-is-the-metaverse/>, 2022c.
- NVIDIA. Nvdl deep learning accelerator. <http://nvidia.org>, 2017.
- Palmero, C., Sharma, A., Behrendt, K., Krishnakumar, K., Komogortsev, O. V., and Talathi, S. S. Openeds2020 challenge on gaze tracking for vr: Dataset and results. *Sensors*, 21(14):4769, 2021.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.
- Qi, C. R., Liu, W., Wu, C., Su, H., and Guibas, L. J. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 918–927, 2018.
- Qualcomm. Snapdragon 888 5g mobile platform. https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/prod_brief_qcom_sd888_5g_0.pdf, 2022.
- Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., and Koltun, V. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- Reddi, V., Kanter, D., Mattson, P., Duke, J., Nguyen, T., Chukka, R., Shiring, K., Tan, K.-S., Charlebois, M., Chou, W., El-Khamy, M., Hong, J., St John, T., Trinh, C., Buch, M., Mazumder, M., Markovic, R., Atta, T., Cakir, F., Charkhabi, M., Chen, X., Chiang, C.-M., Dexter, D., Heo, T., Schmuelling, G., Shabani, M., and Zika, D. Mlperf mobile inference benchmark: An industry-standard open-source machine learning benchmark for on-device ai. In Marculescu, D., Chi, Y., and Wu, C. (eds.), *Proceedings of Machine Learning and Systems*, volume 4, pp. 352–369, 2022. URL <https://proceedings.mlsys.org/paper/2022/file/>

7eabe3a1649ffa2b3ff8c02ebfd5659f-Paper.pdf.

Reddi, V. J., Cheng, C., Kanter, D., Mattson, P., Schmuelling, G., Wu, C.-J., Anderson, B., Breughe, M., Charlebois, M., Chou, W., et al. Mlperf inference benchmark. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 446–459. IEEE, 2020.

Shi, Y., Wang, Y., Wu, C., Yeh, C.-F., Chan, J., Zhang, F., Le, D., and Seltzer, M. Emformer: Efficient memory transformer based acoustic model for low latency streaming speech recognition. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6783–6787. IEEE, 2021.

Tang, R. and Lin, J. Deep residual learning for small-footprint keyword spotting. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5484–5488. IEEE, 2018.

You, H., Wan, C., Zhao, Y., Yu, Z., Fu, Y., Yuan, J., Wu, S., Zhang, S., Zhang, Y., Li, C., et al. Eyecod: eye tracking system acceleration via flatcam-based algorithm & accelerator co-design. *arXiv preprint arXiv:2206.00877*, 2022.

Zhang, J., Jiao, J., Chen, M., Qu, L., Xu, X., and Yang, Q. A hand pose tracking benchmark from stereo matching. In *2017 IEEE International Conference on Image Processing (ICIP)*, pp. 982–986, 2017. doi: 10.1109/ICIP.2017.8296428.

A BENCHMARK MODEL INSTANCES

As an extension to Table 1, this section describes more details on the models included in XRBENCH. Table 1 specifies which model variation (Model Instance) is adopted from the representative model (Model Reference), along with baseline or backbone structure (Model Type), and types of major operators that compose the model (Major Operators). The model instances are chosen based on their size, considering the edge use case. In addition, we also down-scale the dataset resolution of certain tasks to adjust to the context of edge devices. Stereo Hand Pose (Zhang et al., 2017) is scaled by 1/2 for Hand Tracking (HT), OpenEDS 2019 (Garbin et al., 2019) and OpenEDS 2020 (Palmero et al., 2021) are both scaled by 1/4 for Eye Segmentation (ES) and Gaze Estimation (GE), respectively, and KITTI (Geiger et al., 2012) is scaled by 1/4 for Plane Detection (PD).

As shown in the table, there is a variety of model types and operators included in the XRBENCH workloads, representative of the diverse computing requirements of an XR system. Such heterogeneity emphasizes the need for innovative solutions to realize XR device capabilities.

B PROBLEM FORMULATION

We formulate the benchmark and scores using symbols presented in Table 4 and Box 1. We provide more details about the definitions in Box 1 and Box 2 in this section.

Definition 1. Input Data Stream (St_{input})

The input data stream I is defined as follows:

$$St_{input} = \{\sigma \mid \sigma = (inSrcID, FPS_{sensor}, L_{init}, Jt)\}$$

Definition 1 formulates the input stream description in Table 3. The $inSrcID$ refers to a string value that refers to the input source identifier. FPS_{sensor} refers to the streaming rate (FPS) of the associated sensor. L_{init} refers to the initial latency of each input stream, and Jt refers to the maximum jitter in milliseconds.

Definition 2. Model Quality Goal (Q)

$$Q = (QM_{ID}, QM_{Targ}, QM_{Type})$$

Model quality refers to the degree of achieved target metrics (e.g., mIoU and accuracy) of each model. QM_{ID} refers to the name of the metric, QM_{Targ} refers to the float point value representing the target value of the metric (e.g., 0.96 for classification accuracy), and QM_{Type} indicates if the metric is higher or lower-is-better (HiB or LiB) metric.

Definition 3. Unit models (M)

The set of unit models M is defined as follows:

$$M = \{\mu \mid \mu \in (M_{ID}, DS_{ID}, \sigma, Q) \wedge \sigma \in St_{input}\}$$

Definition 3 defines a set of unit models utilized in XRBENCH to construct complex usage scenarios. μ refers to a unit model (i.e., an element of M), M_{ID} refers to the model name in string, DS_{ID} refers to the name of associated dataset, σ refers to an input stream from a sensor associated with the unit model, Q is the model quality goal defined in Definition 2. If a model utilizes multi-modality inputs, σ becomes a set of associated input streams. Based on the definition of M , we define the usage scenario (θ) as follows:

Definition 4. Usage Scenario (θ)

$$\theta = \{(\mu, Dep_{\mu}, FPS_{model}) \mid \mu \in M \wedge Dep_{\mu} \subset M\}$$

In Definition 4, Dep_{μ} defines the model granularity dependency on μ , which is a list of models on which μ depends. With Definition 4, we can define the benchmark suite as follows:

Definition 5. Benchmark Suite (Ω)

Given a set of usage scenarios Θ , a real-time MTMM benchmark suite Ω is defined as follows:

$$\Omega = \{\theta_1, \theta_2, \dots, \theta_{NumScn}\}$$

The Definition 5 shows that a real-time MTMM benchmark is a collection of usage scenarios as described in Table 2. $NumScn$ refers to the number of usage scenarios XRBENCH includes.

Based on the formulation on workload side so far, we define some additional concepts for defining XRBENCH's scoring metrics.

Definition 6. Inference Request (IR)

$$IR = (\mu, InFrame_{ID})$$

Using Definition 3 and Definition 6, we define the inference request time and deadline as follows:

Definition 7. Inference Request time ($T_{req}(IR)$)

$$T_{req}(IR) = L_{init}(inSrcID) + \frac{InFrame_{ID}}{FPS_{Sensor}(inSrcID)}$$

$$+ 2Jt (Dist(rand(inSrcID \times InFrame_{ID})) - 0.5)$$

$$where Dist(x) \in [0, 1] \wedge x \in \mathbb{R}$$

$L_{init}(inSrcID)$ indicates the setup latency of the input stream from the input source $inSrcID$. $InFrame_{ID} \times 1/FPS_{Sensor}(inSrcID)$ represents the time until an XR device reaches the $InFrame_{ID}$ frame under the streaming rate (FPS) of a corresponding input stream σ . The term, $Jt \times (Dist(rand(inSrcID \times InFrame_{ID})) - 0.5)$,

Table 7. Specific model instances for the XRBENCH unit models listed in Table 1. Also classifies model type and major operators.

Task	Model Reference	Model Instance	Model Type	Major Operators
HT	Hand Graph-CNN (Ge et al., 2019)	Hand Shape/Pose (Ge et al., 2019)	CNN	CONV2D, Maxpool, FC
ES	RITNet (Chaudhary et al., 2019)	RITNet (RIT, 2019)	CNN	CONV2D, AvgPool, Skip Connection
GE	EyeCod (You et al., 2022)	FBNet-C (Meta, 2019)	CNN	CONV2D, DWCONV, Skip Connection
KD	Key-Res-15 (Tang & Lin, 2018)	res8-narrow (Tang & Lin, 2018)	CNN	CONV2D, Avgpool, Skip Connection
SR	Emformer (Shi et al., 2021)	EM-24L (Shi et al., 2021)	Transformer	Self-attention, Layernorm
SS	HRViT (Gu et al., 2022)	HRViT-b1 (HRV, 2022)	Transformer	Self-attention, Layernorm, DWCONV
OD	D2go (Meta, 2022a)	Faster-RCNN-FBNetV3A (Meta, 2022a)	R-CNN	CONV2D, DWCONV, Skip Connection
AS	TCN (Lea et al., 2017)	ED-TCN (ED-, 2016)	CNN	CONV2D, Maxpool, Upsample
DE	MiDaS (Ranftl et al., 2020)	midas_v21_small (mid, 2020)	CNN	CONV2D, DWCONV, Skip Connection
DR	Sparse-to-Dense (Ma & Karaman, 2018)	RGBd-200 (Ma & Karaman, 2018)	CNN	CONV2D, DeCONV, DWCONV
PD	PlaneRCNN (Liu et al., 2019)	PlaneRCNN (Liu et al., 2019)	R-CNN	CONV2D w/ FPN, RoIAlign

accounts for the impact of jitters on the arrival time modeled by a maximum jitter Jt , a distribution $Dist(x) \in [0, 1]$ for $x \in \mathbb{R}$, and a random function $rand = f : \mathbb{N} \rightarrow \mathbb{R}$. Note that we make the choice of $Dist(x)$ and $rand(n)$ flexible for various scenarios. By default, $Dist(x)$ is a Gaussian distribution and $rand(n)$ is the rand function of C++17 standard library, *cslib*.

Using a similar formulation, we define the inference deadline as follows.

Definition 8. Inference Deadline ($T_{dl}(IR)$)

The deadline for an inference request IR is defined as follows:

$$T_{dl}(IR) = L_{init}(inSrcID) + \frac{InFrameID + 1}{SR(inSrcID)}$$

The definition in Definition 8 indicates that the deadline of an inference on an input frame is the arrival time of the next input frame.

Definition 9. Inference Slack ($T_{sl}(IR)$)

The inference slack, the length of time window given for an inference run associated with IR is defined as follows:

$$T_{sl}(IR) = T_{dl}(IR) - T_{req}(IR)$$

Using the definitions from Definition 1 to Definition 9, we define XRBENCH’s score metrics.

B.1 Inference-level Benchmark Score

Because of the real-time processing nature of the metaverse workloads, it is challenging to compare metaverse device systems running real-time MTMM using traditional metrics of using latency and energy. The latency measures the end-to-end execution time of each inference, which can be used to check if each model’s deadlines are satisfied. However, achieving less latency than the deadlines does not offer benefits, making latency not an absolute minimization goal, unlike other ML systems targeting non-real-time MTMM workloads. In addition, we can adjust energy to meet the

deadlines or optimize using the slack to the deadline (e.g., DVFS), which also makes energy a knob, not an absolute minimization target.

Therefore, we explore a set of new metrics for XRBENCH that considers all the unique aspects of real-time MTMM workloads we discussed: (1) the task-specific deadlines based on usage scenarios, (2) end-to-end latency (i.e., how much latency does an ML system need beyond the deadline?), (3) overall energy consumption, and (4) the quality of experience delivered.

To facilitate an intuitive comparison of ML systems with many pillars, we propose a single score metric that captures all of the above. The single-score approach will also help motivate the industry to submit their results because the latency and energy can be confidential data that cannot be directly shared with the public. Instead, the industry can offer the single score metric capturing overall performance on real-time and multi-model DL workloads to demonstrate the robust capabilities of their accelerator systems.

To construct such a metric, we consider each aspect of real-time MTMM workloads and model performance (e.g., accuracy, mIoU, and boxAP): real-time requirement, energy consumption, quality of experience, and model performance. We define a score for each pillar and combine them to define the single comprehensive metric. We define those score functions to be in $[0, 1]$ range to facilitate component-wise analysis as well.

To model real-time requirements, we consider the following observations: (1) Too much optimization on latency beyond the deadline does not lead to higher processing rates; even if a system finished the latency in only one cycle, the system still needs to wait for the next input frame. (2) reduced latency can still be helpful for scheduling other models. (3) violated deadline gradually damages the user experience (e.g., Achieving 59Hz for an eye-tracking model targeting 60Hz wouldn’t significantly affect the user experience).

Based on those observations, we search for a function that (1) gradually rewards and penalizes the reduced and increased latency near a deadline (e.g., $\pm 0.5ms$ for a deadline of 10ms) and (2) outputs 0 and 1 if the latency is beyond or




Figure 8. An example real-time score function over different values of the parameter k whose range is $[0, \infty)$. We assume the time window between the inference request time and deadline to be 1 s in this example for simplicity. If k is 0, the score is completely not relevant to the deadline (i.e., no sensitivity on deadline). If k is ∞ , the score function becomes a piece-wise function that flips the score from 1 to 0 at the deadline.

within the deadline. We discuss such a function in Definition 10.

Definition 10. Real-time (RT) score

For an inference request IR , the real-time score is defined as follows:

$$RtScore(IR) = \frac{1}{1 + e^{k(L_{Inf}(IR) - T_{sl}(IR))}}$$

The definition of $RtScore(IR)$ in Definition 10 is based on the inference latency $L_{Inf}(IR)$ and the time window given for the inference IR , $T_{sl}(IR)$.

The RT score function can also be viewed as a modified Sigmoid function with shifting and scaling. A benefit of using $RtScore$ is that we can tune the constant k depending on the deadline sensitivity. Figure 8 shows the change of the $RtScore$ based on the k value, and we can observe large k values makes the function more sensitive around the deadline. We set 15 as the default value of k and utilize it in our evaluation.

Definition 11. Energy (En) score

For an inference request IR , the energy score is defined as follows:

$$EnScore(IR) = \frac{En_{max} - En(IR)}{En_{max}}$$

To make the energy score as higher-is-better metric, consistent with other scores, we utilize En_{max} , which represents the maximum energy allowed for each inference. By default, we set En_{max} as 1500 mJ.

Definition 12. Accuracy (Acc) score

$$AccScore(IR) = \max(1, rawAccScore(IR))$$

$$rawAccScore(IR) = \begin{cases} \frac{QM_{measured}}{QM_{target}}, & \text{if } QM_{Type} = HiB \\ \frac{QM_{measured}}{QM_{measured} + \epsilon}, & \text{otherwise} \end{cases}$$

where $\epsilon > 0 \wedge \epsilon \ll 1 \wedge \epsilon \in \mathbb{R}$

For an inference request IR , the accuracy (Acc) score quantifies the model quality (or model performance) as a value within the $[0,1]$ range. Depending on the model quality type (higher- and lower-is-better), Acc score is computed in a different way to formulate the Acc score as a higher-is-better metric. For lower-is-better model quality metric, we utilize a small real number ϵ for numerical stability, preventing divide by zero errors. By default, ϵ is set as 10^{-6} .

The three base scores, RT, EN, and Acc Scores, are defined for each inference runs. Unlike those, the next score, quality of experience score is defined for entire inference runs for a model.

Definition 13. Quality of Experience (QoE) score

For a unit model μ , the quality of experience (QoE) score is defined as follows:

$$QoEScore(\mu) = \frac{NumFrm_{exec}(\mu)}{NumFrm(\mu)}$$

$NumFrm(\mu)$ refers to the total number of input frames for a unit model μ streamed during the entire execution of the workload. $NumFrm_{exec}(\mu)$ refers to the number of actually processed input frames using μ .

The QoE score quantifies the ratio of the number of processed input frames and the total number of streamed input frames for a model μ . The QoE score is defined in a usage scenario granularity because frame drop can be measured for the entire inference runs for a model.

Using four unit scores (RT, En, Acc, and QoE scores,), we formulate the inference, usage scenario, and benchmark scores.

Definition 14. Inference-wise score

$$Score_{inf}(IR) = RtScore(IR) \times EnScore(IR) \times AccScore(IR)$$

As illustrated in Figure 4, we compute the product of three unit scores defined in the inference granularity to define an aggregated metric for an inference. Combining the inference-wise score with the QoE score, we construct the usage scenario granularity score.

Definition 15. Usage-scenario Score

$$Score_{scn}(\theta) = \sum_{j=1}^{NumFrm(\mu)} \frac{Score_{inf}(IR) \times QoEScore(\mu)}{NumFrm(\mu) \times |\theta|}$$

Using usage-scenario-wise scores, we define an overall score for XRBENCH, XRBENCH SCORE, as the average of the scores for each usage scenario in XRBENCH.

Definition 16. XRBENCH SCORE

$$Score_{bench} = \frac{\sum_{\theta \in \Omega} Score_{scn}(\theta)}{|\Omega|}$$

As shown in Definition 16, XRBENCH SCORE ($Score_{bench}$) summarizes scores for each usage scenario using average.

B.2 Schedule

We do not propose a specific scheduler as it is a part of the ML system software to be evaluated. However we define valid schedules to satisfy the following conditions:

Dependency Condition

$$\forall h_1, h_2 \in HW, \mu_i \in Dep_{\mu_j}, T_{end}(\mu_i, h_1) < T_{start}(\mu_j, h_2)$$

Hardware Occupancy Condition

$$\forall \mu_1, \mu_2 \in M,$$

$$(T_{end}(\mu_1, h) \leq T_{start}(\mu_2, h)) \vee (T_{end}(\mu_2, h) \leq T_{start}(\mu_1, h))$$

where $T_{end}(h, \mu) = \infty \wedge T_{start}(h, \mu) = 0$ if μ is not mapped on h

The dependency condition indicates that the dependency order must be maintained in a schedule. The hardware occupancy condition indicates that a hardware piece (e.g., a systolic array-based accelerator) cannot run two models simultaneously. That is, if a hardware piece can run multiple models simultaneously, it should be treated as multiple smaller hardware pieces.

C DETAILED RELATED WORK COMPARISON

In this section, we expand on Section 5 and Table 6 by providing detailed discussions on prior benchmarks.

General ML Workload Benchmarks. MLPerf Inference (Reddi et al., 2020) is a set of industry standard, single-kernel ML benchmarks that span the ML landscape, from high performance computers (Farrell et al., 2021) to tiny embedded systems (Banbury et al., 2021). It also provides a rich set of inference scenarios based on realistic use cases from industry: single-stream (single inference), multistream (repeated inference with a time interval), server (random inference request modeled via Poisson distribution), and offline (batch processing). Extensions to embedded systems

(MLPerf Tiny (Banbury et al., 2021)) and mobile devices like smartphones (MLPerf Mobile (Reddi et al., 2022)) have also been developed, drawing closer to the XR form factor. However, the MLPerf suite workloads do not deploy models in a concurrent or cascaded manner and the scoring metrics lack QoE consideration, which are essential in XR workloads.

DeepBench (Dee, 2016) focuses on benchmarking kernel operations which underlie ML performance. Although such microbenchmarks provide insights to operator level optimizations, it cannot be used for understanding the end-to-end performance of a single model or for MTMM workloads. AI Benchmark (Ignatov et al., 2018) targets the ML inference performance of smartphones with 14 different tasks and EEMBC MLMark (EEM, 2020) measures the performance of neural networks on embedded devices. Still, none of them cover MTMM performance nor consider real-time processing scenarios. Their scoring metrics are also not sufficiently diverse to handle complex XR workloads.

AIBench (Gao et al., 2019) from BenchCouncil is another industry standard AI benchmark for Internet services, which was one of the first to include application scenarios for end-to-end performance evaluation. These scenarios model MTMM workloads of E-commerce search intelligence use cases with heterogeneous latency of each model, provided with rich scoring metric components for evaluation. Although AIBench decently reflects the key components of real-time MTMM workloads, the benchmark is tailored to server-scale internet service and has little to do with edge applications. In addition, their static execution graphs make extensions to XR use cases difficult, which require dynamic execution of models based on their control dependencies.

AIoT (Aio, 2018; Luo et al., 2018) is an AIBench extension that focuses on mobile and embedded AI. Though these platforms come closer to the XR platform, the benchmark does not model real-time, MTMM-based scenarios and therefore falls short to serve as an XR benchmark.

XR Benchmarks. ILLIXR (Huzaifa et al., 2021) is a benchmark suite tailored for XR systems. ILLIXR models concurrent and cascaded execution pipelines in XR use cases and considers the real-time requirements of XR devices. Although ILLIXR provides a solid benchmark in the XR domain, the focus of ILLIXR is mainly in non-ML-based pipelines, unlike the ML workload focus of XRBENCH. ILLIXR includes one ML model (RITNet for eye tracking), and its other parts are based on traditional computer vision and audio algorithms (e.g., QR decomposition and Gauss-Newton refinement) and signal processing (e.g., FFT).

VRMark (VRM, 2020) is a benchmark that evaluates the performance of VR experiences on PCs. The benchmark also does not target ML performance assessments but rather

focuses on rendering graphics. Moreover, it lacks usage scenarios that are reflective of real-world user characteristics and various score metrics for systematical analysis.

ML and XR Benchmarks. Compared to the above-mentioned benchmarks (Dee, 2016; EEM, 2020; VRM, 2020; Banbury et al., 2021; Farrell et al., 2021; Gao et al., 2019; Huzaifa et al., 2021; Ignatov et al., 2018; Luo et al., 2018; Reddi et al., 2020), XRBENCH covers all requirement of an ML-based XR workloads. To be specific, XRBENCH provides diverse cascon-MTMM scenarios with real-time requirement and complex dependencies, which majority ML benchmarks are missing. Careful consideration of QoE aspects in XR applications into its scoring metric is another strength of XRBENCH that distinguishes it from other prior works.

On the other hand, even though existing XR-related or scenario-based benchmarks support real-time MTMM scenario and QoE metrics, they still lack several components such as sufficient ML algorithm coverage, dynamic model execution graph, and focusing on edge devices. All of these characteristics are satisfied by XRBENCH, expecting significant contribution to XR research community and the industry.

D ARTIFACT APPENDIX

D.1 Abstract

This appendix describes the complete workflow for running XRBench-MAESTRO and generating results used in the paper.

D.2 Artifact check-list (meta-information)

- **Algorithm:** Scheduling is based on a latency-greedy scheduler; dispatch an inference job to an idle accelerator with the minimal expected latency.
- **Program:** Based on MAESTRO (<https://maestro.ece.gatech.edu>)
- **Compilation:** C++ compilers that support C++17 or later (tested compilers: clang-1400.0.29.202 or g++ (Ubuntu 11.3.0-1ubuntu1 22.04))
- **Run-time environment:** Tested environments: MacOS 13.2.1 (22D68) and Ubuntu 22.04
- **Hardware:** X86-64 processor-based Linux machines, X86-64 processor-based Mac Machines (e.g., MacBook Pro and iMac), or Apple Silicon-based Mac Machines (e.g., MacBook pro with M1 processor)
- **Execution:** Automated scripts are included (please refer to README in the code for details)
- **Metrics:** score metrics proposed in the paper
- **Output:** Plots in pdf files (under /XR-bench_evaluation/plots), data in csv files (under /XR-bench_evaluation/eval_data)

- **How much disk space required (approximately)?:** 10-20 MB
- **How much time is needed to prepare workflow (approximately)?:** expected to be less than 30 minutes.
- **How much time is needed to complete experiments (approximately)?:** Depends on the machine’s computing power. On our tested machines with latest processors (e.g., Intel i9-13900k with 128GB RAM and Apple M1 Pro with 64 GB RAM), the experiments overall take less than 30 minutes.
- **Publicly available?:** doi: [10.5281/zenodo.7857382](https://doi.org/10.5281/zenodo.7857382), Github: <https://github.com/XRBench>

D.3 Description

D.3.1 How delivered

A Dropbox download link of a zip file of the code will be shared with the AE reviewers (please see “abstract” in the Hotcrp submission). We will open-source a clean version along with the conference.

D.3.2 Hardware dependencies

Typical desktop, laptop, or server running Linux or Mac OS are required. We tested the artifact on X86-64 processor-based Linux machines, X86-64 processor-based Mac Machines (e.g., MacBook Pro and iMac), and Apple Silicon-based Mac Machines (e.g., Macbook pro with M1 processor)

D.3.3 Software dependencies

- C++ compiler supporting C++17 (tested compilers: clang and g++)
- scon
- boost library (C++)
- Python3 (3.10 or later)
- matplotlib

D.3.4 Data sets

D.4 Installation

SW dependency installation guide for Linux machines (Ubuntu)

```
> sudo apt-get install g++ libboost-all-dev \
python3 scon python3-pip
> pip3 install matplotlib
```

SW dependency installation guide Guide for Mac machines (Using Homebrew)

```
> brew install scon python3 boost
> pip3 install matplotlib
```

Note: Apple silicon-based Mac machines have issues in linking boost library during compilation. An alternative SConstruct file (specification file for scon-based compilation flow) that addresses this issue is included in the codebase; please refer to README file for details.)

D.5 Experiment workflow

```
> scons
> ./reproduce_figure5.sh
> ./reproduce_figure6.sh
> ./reproduce_figure7_data.sh
```

D.6 Evaluation and expected result

Reproducing Figure 5 Running “reproduce_figure5.sh,” XRBench-MAESTRO will reproduce plots in Figure 5 of the original paper under “XRbench_evaluation/plots/4K” and “XRbench_evaluation/plots/8K.” Please note that results on outdoor activity A, outdoor activity B, and AR assistant are non-deterministic (the workload is dynamic). Due to such dynamic workloads, “gross_average” plots will look slightly different as well. You can still expect exact match of results on social interaction A, social interaction B, and AR gaming.

Reproducing Figure 6 Running “reproduce_figure6.sh” after “reproduce_figure5.sh,” XRBench-MAESTRO will reproduce plots in Figure 6 of the original paper under “XRbench_evaluation/plots.” You can expect exact match of the results. Please note that the aspect ratio of the figure 6 and generated plots are different (and Figure 6 has mis-aligned x-axis tick labels), but they will show the same data.

Generating Figure 7 data Running “reproduce_figure7_data.sh” after “reproduce_figure5.sh,” and “reproduce_figure6.sh,” data under Figure 7 will be generated under “XRbench_evaluation/eval_data/figure7/.”

Please expect some fluctuations in results as the workload for Figure 7 is dynamic.

D.7 Experiment customization

The provided automated scripts cover entire evaluation we ran. Users can change settings by modifying the contents of files under the followings.

- **Modifying dataflow styles of accelerators:** Please modify “XRbench_evaluation/dataflows” to change processing style of accelerators (i.e, dataflow). The description is based on MAESTRO (<https://maestro.ece.gatech.edu>) style dataflow notation.
- **Modifying hardware styles:** Please modify “XRbench_evaluation/hw_configs” to change hardware parameters (e.g., number of PEs, number of sub-accelerators in the hardware system, etc.)