

# Cornserve: A Distributed Serving System for Any-to-Any Multimodal Models

Jae-Won Chung<sup>1,\*</sup> Jeff J. Ma<sup>1,\*</sup> Jisang Ahn<sup>1</sup> Yizhuo Liang<sup>2</sup>  
Akshay Jajoo<sup>3</sup> Myungjin Lee<sup>3</sup> Mosharaf Chowdhury<sup>1</sup>

<sup>1</sup>University of Michigan <sup>2</sup>University of Southern California <sup>3</sup>Cisco Research

## Abstract

Any-to-Any models are an emerging class of multimodal models that accept combinations of multimodal data (e.g., text, image, video, audio) as input and generate them as output. Serving these models are challenging; different requests with different input and output modalities traverse different paths through the model computation graph, and each component of the model have different scaling characteristics.

We present Cornserve, a distributed serving system for generic Any-to-Any models. Cornserve provides a flexible task abstraction for expressing Any-to-Any model computation graphs, enabling component disaggregation and independent scaling. The distributed runtime dispatches compute to the data plane via an efficient record-and-replay execution model that keeps track of data dependencies, and forwards tensor data between components directly from the producer to the consumer. Built on Kubernetes with approximately 23K new lines of Python, Cornserve supports diverse Any-to-Any models and delivers up to 3.81× higher throughput and 5.79× lower tail latency. Cornserve is open-source,<sup>1</sup> and the demo video is available on YouTube.<sup>2</sup>

## CCS Concepts

• **Computer systems organization** → **Distributed architectures**; • **Computing methodologies** → *Machine learning*.

## Keywords

Multimodal models, Any-to-Any models, Inference serving, Distributed systems

## ACM Reference Format:

Jae-Won Chung, Jeff J. Ma, Jisang Ahn, Yizhuo Liang, Akshay Jajoo, Myungjin Lee, and Mosharaf Chowdhury. 2026. Cornserve: A Distributed Serving System for Any-to-Any Multimodal Models.

\*Equal contribution.

<sup>1</sup><https://github.com/cornserve-ai/cornserve>

<sup>2</sup><https://www.youtube.com/watch?v=nb8R-vztLRg>



This work is licensed under a Creative Commons Attribution 4.0 International License.

CAIS '26, San Jose, CA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2415-2/2026/05

<https://doi.org/10.1145/3786335.3813209>

Model	Input	Output
Qwen 2.5 Omni [18], Qwen 3 Omni [19]	T, I, V, A	T, A
Qwen 2.5 VL [2], InternVL 3 [23]	T, I, V	T
DeepSeek Janus [3, 16]	T, I	T, I
LTX-2 [5]	T, I	V, A
Qwen Image [15], GLM Image [21]	T	I

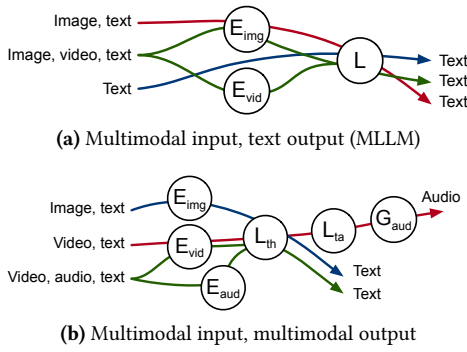
**Table 1:** Modality breakdown of recent Any-to-Any multimodal models. The input and output modalities (Text, Image, Video, Audio) supported by a model can vary significantly.

In *ACM Conference on AI and Agent Systems (CAIS '26)*, May 26–29, 2026, San Jose, CA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3786335.3813209>

## 1 Introduction

Our world is inherently multimodal, abound with data in the form of text, image, video, audio, and more. Naturally, recent model developments generalize text-only Large Language Models (LLMs) and embrace *multimodality*, leading to the emergence of a new class of models called *Any-to-Any* models. With over 11,000 variants on Hugging Face Hub as of March 2026 [6], Any-to-Any models can (1) understand multimodal inputs and/or (2) generate multimodal outputs, with text being the most common modality. For instance, Multimodal LLMs (MLLMs) like Qwen 2.5 VL [2] and InternVL 3 [23] can process multimodal inputs and generate text responses; Qwen Image [10, 15] and GLM Image [21] produce images from text inputs processed by an LLM; LTX-2 [5] generates audio and video; DeepSeek Janus [3, 16] both understands and generates text and images; and Qwen Omni [18, 19] takes text, image, video, and audio inputs and generates text and audio. Table 1 summarizes modalities supported by recent Any-to-Any models.

Any-to-Any models can be viewed as a *graph of heterogeneous components* that process different modalities, such as multimodal encoders, one or more LLMs, and multimodal generators (Figure 1). This introduces two types of heterogeneity to the serving system. First is request type and computation path heterogeneity. For instance, the Qwen Omni model in Figure 1b consists of an image, video, and audio encoders that feed multimodal embeddings into the *thinker* LLM for text generation. If audio output is requested by the user, the text output and hidden states of the thinker LLM is passed to another autoregressive component called the *talker* that outputs audio tokens, which are subsequently decoded into waveforms by an audio generator (vocoder). Thus,



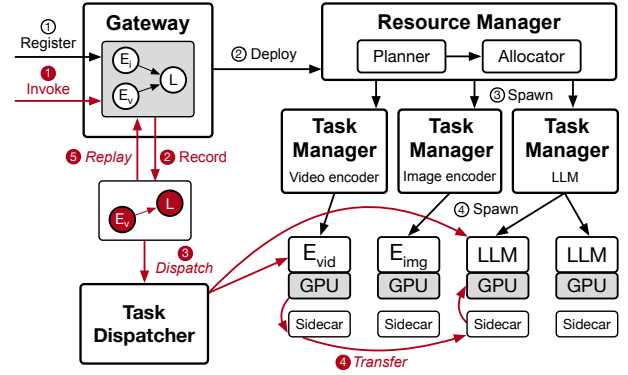
**Figure 1:** Computation graphs of (a) InternVL 3 [23], a multimodal input model, and (b) Qwen Omni [18, 19], a multimodal input and output model. Different requests invoke different components and take different paths on the graph.  $E$  stands for Encoder,  $L$  for LLM, and  $G$  for Generator.  $L_{th}$  and  $L_{ta}$  stand for thinker and talker LLMs, respectively.

inference requests with different combinations of input and output modalities lead to different paths through the graph and invoke different model components, resulting in different request rates for each component. This is exacerbated by the second type of heterogeneity: different components have different scaling characteristics and resource requirements. For example, in Qwen 2.5 Omni, the text generation component achieves 4× higher throughput than the audio generator on the same A100 GPU.

Existing serving systems have offered at best *point solutions* for Any-to-Any models. Systems like vLLM [7] have focused on text-only LLMs and MLLMs, whereas xDiT [4] has focused only on diffusion-based image or video generation, i.e., special cases of Any-to-Any models. Similarly, techniques like Prefill–Decode (PD) disaggregation [9, 22] and Encode–Prefill–Decode (EPD) disaggregation [11] target specific model architectures and serving scenarios.

To this end, we present Cornserve, the first distributed serving system for generic Any-to-Any models, to the best of our knowledge. Cornserve provides three key capabilities for efficiently serving Any-to-Any models:

- A flexible **task abstraction** (unit tasks, composite tasks, and apps) that allows model developers to express arbitrary Any-to-Any model computations in plain Python.
- **Model fission** to disaggregate models at component boundaries into independently scalable components, each running on dedicated GPUs with specialized executors.
- A **distributed runtime** that supports models with arbitrary computation graphs through a record-and-replay execution mechanism, with efficient tensor data forwarding between components via shared memory and RDMA.



**Figure 2:** Cornserve architecture and flows for deployment (black) and inference (red).

Additionally, when multiple models use equivalent components (e.g., a shared vision encoder), Cornserve automatically shares executor deployments to reduce GPU usage.

Cornserve is implemented on top of Kubernetes with approximately 23K lines of Python, with full support for OpenTelemetry for monitoring and tracing. Cornserve supports a variety of recent Any-to-Any models including Qwen 2.5 VL [2], InternVL 3 [23], Qwen Image [15], and Qwen 2.5/3 Omni [19], and improves serving throughput by up to 3.81× and reduces tail latency by up to 5.79×.

## 2 System Design

### 2.1 Architecture Overview

Figure 2 illustrates Cornserve’s overall architecture and how deployment and inference requests flow through the system.

**Control Plane.** The *Gateway* serves as the entry point for *App* registration and inference requests. An *App* is a user-defined Python module that specifies one or more *Tasks*, each of which specifies a model component and its dependencies on other components. The *Resource Manager* includes a *planner* [8] that finds an efficient deployment plan for tasks, and an *allocator* that realizes the plan by spawning *Task Managers*. Task Managers (1) spawn *Task Executors* (e.g., vLLM) on the allocated GPUs and (2) handles routing/load balancing. The *Task Dispatcher* routes component invocations from the Gateway to the appropriate Task Executors.

**Data Plane.** Each *Task Executor* runs model components on GPUs. Cornserve provides specialized executors for different component types: *Eric* for multimodal encoders (image, video, audio), a forked vLLM [7] for LLM inference, and *Geri* for multimodal generators (audio vocoder, image DiT). The *Sidecar*, which is a per-GPU daemon process, forwards tensor data between executors following data dependencies.

**Deployment Lifecycle.** An administrator (1) registers an app by submitting a Python source file to the Gateway. The

Gateway validates the app, identifies its task definitions, and ② invokes the Resource Manager. The Resource Manager ③ spawns Task Managers, which in turn ④ spawn Task Executor replicas on the allocated GPUs.

**Request Lifecycle.** When ① an inference request arrives at the Gateway, it ② identifies the specific components invoked by the request using a *record-and-replay* approach (§2.2) and ③ forwards them to the Task Dispatcher, which dispatches these to the Task Executors. Intermediate data ④ flows through Sidecars, bypassing the control plane, and the final result is ⑤ returned to the client. All component invocations are dispatched simultaneously, and Task Executors wait for intermediate tensors to arrive through the Sidecar.

## 2.2 Task Abstraction and Model Fission

Any-to-Any models exhibit request and computation path heterogeneity: depending on the request type, the model may invoke different components in different sequences. For instance, in a multimodal LLM with an image encoder and an LLM, a request with both image and text invokes the encoder first and then the LLM, while a text-only request bypasses the encoder and directly invokes the LLM. In the general case, this logic becomes arbitrarily complex with loops and branches that depend on the input, and the inference logic may include arbitrary processing before and after invoking a component (e.g., filtering returned text, modifying prompts).

**Tasks and Apps.** Cornserve introduces three levels of abstraction: *unit tasks*, *composite tasks*, and *apps*. A unit task, e.g., `LLMTask`, `EncoderTask`, `GeneratorTask`, represents an atomic model component, and is deployed as its own Task Manager. Model developers compose unit tasks within a composite task class that implements a model’s computations in plain Python inside the `invoke` method. An *app* is a single Python module that instantiates unit/composite tasks and defines an `async def serve` function as the entry point.

**Model Fission.** By expressing a model as a composite task of unit tasks, users can define *fissions* at arbitrary component boundaries. Listing 1 shows `MLLMTask`, a built-in composite task for multimodal LLMs. When `encoder_fission` is true, `post_init` creates separate `EncoderTask` unit tasks that are deployed on dedicated GPUs, and the `invoke` method calls them to produce embeddings before invoking the LLM. When false, the LLM handles encoding internally in a monolithic deployment. Each unit task is associated with a *task descriptor* that specifies how to deploy the Task Executor for that task.

**Record-and-Replay Execution.** During runtime, for a given inference request, Cornserve must extract the subgraph of unit task invocations and dispatch them to the appropriate Task Executors. The key challenge is that the invocation

```
cornserve_tasklib/task/composite/llm.py
class MLLMTask(Task[ChatCompletion, Stream[str]]):
    model_id: str
    modalities: list[Modality]
    encoder_ids: set[str]
    encoder_fission: bool = True

    def post_init(self):
        if self.encoder_fission:
            self.encoders = {
                mod: EncoderTask(self.encoder_ids, modality)
                for modality in self.modalities
            }
        self.llm = LLMTask(
            self.model_id, recv_embeds=self.encoder_fission)

    def invoke(self, req: ChatCompletion) -> Stream[str]:
        if self.encoder_fission:
            for mes in extract_multimodal(req.messages):
                emb = self.encoders[mes.modality].invoke(mes)
                req.multimodal_embeddings.append(emb)
        return self.llm.invoke(req)
```

**Listing 1:** Simplified `MLLMTask` composite task. When `encoder_fission` is true, `LLMTask` is instantiated with `recv_embeds=True`, which is recognized by its task descriptor to disable the encoder and receive embeddings from the sidecar instead. Cornserve’s *record-and-replay* captures unit task invocations and data dependencies automatically.

```
examples/gemma_arena.py
gemmas = ["google/gemma-3-4b-it", "google/gemma-3-12b-it"]
# All MLLMTasks share the same encoder deployment.
gemma_tasks = {
    model_id: MLLMTask(
        modalities=[Modality.IMAGE],
        model_id=model_id,
        encoder_ids=set(gemmas),
        encoder_fission=True,
    )
    for model_id in gemmas
}

class Config(AppConfig):
    tasks = {**gemma_tasks}

async def serve(req: ChatCompletion) -> AsyncIterator[dict]:
    streams = await asyncio.gather(
        *(task(req) for task in gemma_tasks.values())
    )
    async for model, chunk in merge_streams(gemmas, streams):
        yield {model: chunk}
```

**Listing 2:** Cornserve app serving multiple Gemma 3 models [13]. The `serve` function defines the app’s computation: it invokes multiple `MLLM` tasks concurrently, merges their streaming responses, and yields structured results, demonstrating how apps can express arbitrary logic in plain Python.

subgraph depends on the request content and may involve complex control flow, as illustrated in Listing 1.

Cornserve addresses this with a two-phase approach. In the *record* phase, the composite task’s `invoke` method is

called with the request, but unit tasks return placeholder results while recording their invocations and inputs. By matching placeholder object IDs across unit task inputs and outputs, Cornserve reconstructs the subgraph of the computation graph executed for the given request. This phase is instantaneous, as no actual model computation is performed. The recorded graph is then sent to the Task Dispatcher, which dispatches each unit task invocation to the appropriate executors and collects results. In the *replay* phase, the *invoke* method is called again, but now each unit task returns the real result from the dispatcher, producing the final response.

This record-and-replay approach allows Cornserve to support Any-to-Any models with most control flow constructs (e.g., loops, branches) while keeping the developer experience simple. A limitation is that the composite task’s control flow must be deterministic given the request, as *replay* must follow the same path as the *record*. Dynamic data-dependent control flow (e.g., invoking different tasks based on an LLM’s response) is handled at the app layer, where the *serve* function executes with real results and can have arbitrary logic.


**Component Sharing.** When multiple apps use equivalent unit tasks—same class and field values—they automatically share a single Task Manager and its executors. Listing 2 shows a Cornserve app that serves multiple Gemma 3 models with a shared vision encoder. Because all `MLLMTask` instances specify the same `encoder_ids`, Cornserve deploys only one shared encoder Task Manager, reducing total GPU usage.

**Data Forwarding.** When model components are fissioned onto separate GPUs, intermediate data such as multimodal embeddings and hidden states must be transferred between executors via the Sidecar. For intra-node transfers, Sidecars use Linux shared memory (`/dev/shm`) to minimize interference with NVLink/NVSwitch bandwidth, which is likely heavily utilized by executors running tensor or expert parallel inference [12]. The sender writes intermediate data from GPU memory directly to shared memory and notifies the receiver of the data’s address via gRPC. Upon notification, the receiver reads the data directly from shared memory and copies it into its GPU memory. For inter-node transfers, Sidecars use RDMA via the UCX communication library [1].

### 3 Evaluation

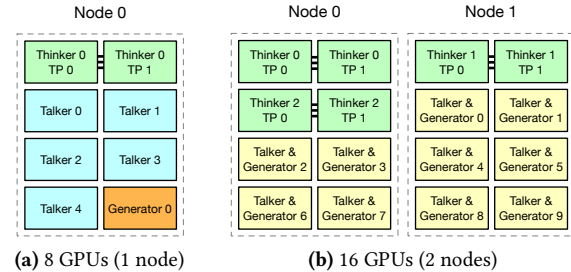
Cornserve supports generic Any-to-Any models including MLLMs. In this section, we highlight Qwen Omni models [18, 19] as a generic and challenging model family to serve.

**Setup.** We conduct experiments on two  $8 \times$  NVIDIA A100-80GB GPU nodes with NVSwitch and 400 Gbps cross-node bandwidth over RDMA. We use ServeGen [17], a real-world multimodal request dataset, and have 20% requests generate audio outputs. The baseline is a monolithic deployment



(a) Qwen 2.5 Omni Tput (b) Qwen 3 Omni Tput (c) Qwen 2.5 Omni Lat.

**Figure 3:** Monolith vs. Cornserve comparisons for Qwen 2.5 Omni 7B [18] throughput and latency CDF, and Qwen 3 Omni 30B [19] throughput.  $\times$  indicate GPU out-of-memory.



**Figure 4:** Cornserve deployment configurations for Qwen 3 Omni on 8 and 16 GPUs. Each box represents a GPU. Model fission allows each component to scale independently: the thinker (LLM) uses tensor parallelism while talkers and generators are replicated to balance throughput.

where all components run within a single executor backed by Hugging Face Transformers [14], as vLLM [7] does not support either Qwen Omni models fully.

**Throughput.** Figures 3a and 3b compare serving Qwen 2.5 Omni and Qwen 3 Omni with monolith vs. Cornserve. Cornserve improves the throughput of Qwen 2.5 Omni on 8-GPU and 16-GPU configurations by  $3.09\times$  and  $3.81\times$ , respectively. Due to the high computation scaling heterogeneity between the LLM and the audio generator, the planner [8] replicates the audio generator independently (7 and 15 replicas on 8 and 16 GPUs, respectively) to balance throughput.

For Qwen 3 Omni, the monolithic deployment fails due to GPU out-of-memory errors, further highlighting the need for model fission. Cornserve’s planner [8] disaggregates Qwen 3 Omni into separately scalable components, enabling efficient serving. Figure 4 shows the resulting deployment configurations: on 8 GPUs, the thinker uses tensor parallelism across 2 GPUs while the talker is replicated across the remaining GPUs with one generator replica; on 16 GPUs, the thinker scales to 3 TP-2 replicas with 10 talker-generator replicas. Scaling from 8 to 16 GPUs yields a  $2.68\times$  throughput improvement, exceeding the expected  $2\times$  from linear scaling as the additional GPUs allow better throughput balancing across components.

**Latency.** As shown in Figure 3c, compared to monolithic deployment, Cornserve significantly reduces request latency for Qwen 2.5 Omni on 16 GPUs, with P50, P95, and P99 improvements of 3.24×, 5.3×, and 5.79×, respectively. These benefits stem from two sources. First, in monolithic deployments, only one component of the model runs at a time, meaning whenever one component is running, all other components do not make progress. Thus, removing this component interference improves latency by reducing blocking. Second, disaggregating the audio generator allows the autoregressive audio generator to unlock continuous batching [20], improving its throughput compared to the monolithic deployment; implementing talker continuous batching alongside the thinker LLM is technically challenging.

**Sidecar Overhead.** Sidecar transfers are not the main latency bottleneck as they only happen at disaggregated component boundaries. Cross-node tensor forwarding via RDMA adds 5–10 ms for 8 MB, 8–20 ms for 16 MB, and 12–27 ms for 32 MB tensors at 5–15 transfers per second; intra-node forwarding via shared memory adds similar overhead.

## 4 Conclusion

We present Cornserve, the first distributed serving system for generic Any-to-Any multimodal models. Cornserve’s task abstraction and model fission enable flexible and efficient deployment of complex Any-to-Any models. We believe the strong inference serving platform provided by Cornserve will enable the development of more powerful and efficient Any-to-Any models, and we hope it will serve as a foundation for future research in this area.

## Acknowledgments

We thank the CAIS reviewers and members of SymbioticLab for their helpful discussions and feedback. This work was supported in part by NSF grants CCF-2450085 and CNS-2106184, and by grants from Cisco, Ford, Mozilla Foundation, and Laude Institute. Jae-Won Chung is additionally supported by the Kwanjeong Educational Foundation.

## References

- [1] The Unified Communication X Library. <http://www.openucx.org>.
- [2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [3] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025.
- [4] Jiarui Fang, Jinzhe Pan, Xibo Sun, Aoyu Li, and Jiannan Wang. xDiT: an inference engine for diffusion transformers (DiTs) with massive parallelism. *arXiv preprint arXiv:2411.01738*, 2024.
- [5] Yoav HaCohen, Benny Brazowski, Nisan Chiprut, Yaki Bitterman, Andrew Kvochko, Avishai Berkowitz, Daniel Shalem, Daphna Lifschitz, Dudu Moshe, Eitan Porat, Eitan Richardson, Guy Shiran, Itay Chachy, Jonathan Chetboun, Michael Finkelson, Michael Kupchick, Nir Zabari, Nitzan Guetta, Noa Kotler, Ofir Bibi, Ori Gordon, Poriya Panet, Roi Benita, Shahar Armon, Victor Kulikov, Yaron Inger, Yonatan Shifitan, Zeev Melumian, and Zeev Farbman. LTX-2: Efficient joint audio-visual foundation model. *arXiv preprint arXiv:2601.03233*, 2026.
- [6] Hugging Face. Any-to-any models on hugging face. [https://huggingface.co/models?pipeline\\_tag=any-to-any](https://huggingface.co/models?pipeline_tag=any-to-any), 2026. Accessed: 2026-03-06.
- [7] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with PagedAttention. In *SOSP*, 2023.
- [8] Jeff J. Ma, Jae-Won Chung, Jisang Ahn, Yizhuo Liang, Runyu Lu, Akshay Jajoo, Myungjin Lee, and Mosharaf Chowdhury. Cornfigurator: Automated planning for any-to-any multimodal model serving. *arXiv preprint arXiv:2512.14098*, 2025.
- [9] Pratyush Patel, Esha Choukse, Chaojie Zhang, Aashaka Shah, İniigo Gori, Saeed Maleki, and Ricardo Bianchini. Splitwise: Efficient generative llm inference using phase splitting. In *ISCA*, 2024.
- [10] Qwen Team. Qwen-Image 2.0. <https://qwen.ai/blog?id=qwen-image-2.0>, 2025.
- [11] Gursimran Singh, Xinglu Wang, Yifan Hu, Timothy Tin Long Yu, Linzi Xing, Wei Jiang, Zhefeng Wang, Bai Xiaolong, Yi Li, Ying Xiong, Yong Zhang, and Zhenan Fan. Efficiently serving large multimodal models using EPD disaggregation. In *ICML*, 2025.
- [12] Biao Sun, Ziming Huang, Hanyu Zhao, Wencong Xiao, Xinyi Zhang, Yong Li, and Wei Lin. Llumix: Dynamic scheduling for large language model serving. In *OSDI*, 2024.
- [13] Gemma Team. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.
- [14] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *EMNLP*, 2020.
- [15] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, Yuxiang Chen, Zecheng Tang, Zekai Zhang, Zhengyi Wang, An Yang, Bowen Yu, Chen Cheng, Dayiheng Liu, Deqing Li, Hang Zhang, Hao Meng, Hu Wei, Jingyuan Ni, Kai Chen, Kuan Cao, Liang Peng, Lin Qu, Minggang Wu, Peng Wang, Shuting Yu, Tingkun Wen, Wensen Feng, Xiaoxiao Xu, Yi Wang, Yichang Zhang, Yongqiang Zhu, Yujia Wu, Yuxuan Cai, and Zenan Liu. Qwen-Image technical report. *arXiv preprint arXiv:2508.02324*, 2025.
- [16] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. *arXiv preprint arXiv:2410.13848*, 2024.
- [17] Yuxing Xiang, Xue Li, Kun Qian, Yan Zhang, Wenyan Yu, Ennan Zhai, Xin Jin, and Jingren Zhou. ServeGen: Workload characterization and generation of large language model serving in production. In *NSDI*, 2026.
- [18] Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-Omni technical report. *arXiv preprint arXiv:2503.20215*, 2025.
- [19] Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfu Zhu, Yuanjun Lv, Yongqi Wang, Dake Guo, He Wang, Linhan Ma, Pei Zhang, Xinyu Zhang, Hongkun Hao, Zishan Guo, Baosong Yang, Bin Zhang, Ziyang Ma, Xipin Wei, Shuai Bai, Keqin Chen, Xuejing Liu, Peng Wang, Mingkun Yang, Dayiheng Liu, Xingzhang Ren, Bo Zheng, Rui Men, Fan Zhou, Bowen Yu, Jianxin Yang, Le Yu, Jingren Zhou, and Junyang Lin. Qwen3-Omni technical report. *arXiv preprint arXiv:2509.17765*, 2025.
- [20] Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. Orca: A distributed serving system for Transformer-Based generative models. In *OSDI*, 2022.
- [21] Z.AI Team. GLM-Image. <https://z.ai/blog/glm-image>, 2025.
- [22] Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. DistServe: Disaggregating prefill and decoding for goodput-optimized large language model serving. In *OSDI*, 2024.
- [23] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, Hao Li, Jiahao Wang, Nianchen Deng, Songze Li, Yanan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingtong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Han Lv, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. InternVL3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.