

A MINDSPEED-MLLM

A.1 ASCEND V.S. NVIDIA

Table 9: Comparison Between NVIDIA CUDA and Ascend CANN.

Component	CUDA	CANN
GE Graph Engine	TensorRT plugins & parser	Graph Engine
Collective Communication Library	NV NCCL	HCCL
Library/Template	NV CUTLASS	Ascend C High-Level API
General Programming	NV CUDA-C	Ascend C Low-Level API
Operator Acceleration Library	NV cuDNN	Ascend aclNN
Runtime	NV Runtime	Ascend Runtime
Driver	NV Driver	Ascend Driver

The ecosystem serves as the cornerstone of distributed pre-training frameworks, with NVIDIA’s CUDA and Huawei’s Ascend CANN being the two mainstream ecosystems. While both possess dedicated driver and runtime layers for OS-AI accelerator communication as illustrated in Table 9, they differ significantly in middle-level acceleration libraries, upper-level tools, and ecological maturity. NVIDIA’s cuDNN and NCCL are highly optimized, and with CUTLASS and TensorRT providing flexible programming templates and graph optimization, it has formed a mature “hardware-software-tools-community” loop, supported by in-depth adaptation of open source communities like those behind PyTorch and Megatron-LM.

In contrast, Ascend’s aclNN, HCCL, and two-tier Ascend C API system have realized core functions but lack coverage of long-tail operators, third-party tools, and community resources. This gap limits Ascend’s training framework selection—mainstream CUDA-based frameworks cannot directly leverage Ascend’s hardware capabilities, and simple adaptation incurs high costs and performance losses, highlighting the necessity of developing a dedicated framework for Ascend to integrate with CANN components seamlessly.

Table 10: Comparison of Differences Between Huawei Ascend and NVIDIA Computing Cards.

XPU	A100	A800	H800	910B1	910B2
TF32 (Tensor Core)	156	156	495	200	188
FP32 (Tensor Core)	NA	NA	NA	100	94
BF16 (Tensor Core)	312	312	989	400	376
FP16 (Tensor Core)	312	312	989	400	376
Int8 (Tensor Core)	624	624	1979	800	752
HBM (GB)	80	80	80	64	64
HBM Bandwidth (GB/s)	2039	2039	3350	1800	1800
NVLink Bandwidth (GB/s)	600	400	600	NA	NA
PCIe Bandwidth (GB/s)	64	64	128	64	64

Beyond underlying chip architecture differences, Ascend and NVIDIA computing cards vary distinctly in core specifications that impact large-model pre-training efficiency and stability as illustrated in 10. In terms of computing power, Ascend 910B shows competitiveness at mainstream precisions and unique advantages in FP32 Tensor Core support. However, Ascend has obvious shortcomings: 64 GB HBM risks training interruptions for ultra-large models, 1800 GB/s HBM bandwidth is lower than NVIDIA’s, and without NVLink, it relies on 64 GB/s PCIe for multi-card communication, increasing latency.

To mitigate these, hardware-friendly strategies are needed, which further drives Ascend-specific framework optimizations—such as developing fused operators, dynamic scheduling, and communication optimization to offset shortcomings and leverage its computing power.

A.2 PRECISION ALIGNMENT

To verify the forward and backward precision of the developed framework, we conducted comparisons not only with the original MindSpeed-MM but also with the following mainstream training frameworks within the NVIDIA ecosystem:

Llama-FactoryZheng et al. (2024): an easy-to-use open-source training framework. It is mainly built on native PyTorch and the HuggingFace ecosystem, and supports distributed engines such as DeepSpeed and FSDP. Currently, its support for the Ascend ecosystem is becoming increasingly comprehensive.

Pai-Megatron-PatchCloud (Year): an open-source training tool developed by Alibaba. Similar to MindSpeed-LLM/MM, it includes adaptation and optimization based on Megatron for mainstream open-source LLMs.




Figure 5: Comparison of loss values among various training tools and platforms on the COCO dataset.

Two datasets were employed to validate the accuracy of the training framework: the public COCO dataset(Lin et al., 2014) and an in-house Slow Thinking dataset. In all training experiments, consistent training hyper-parameters were maintained. Specifically, the training length was set to 4K for the COCO dataset, whereas a 12K training length was adopted for the Slow Thinking dataset.

Firstly, we checked the forward precision of the MindSpeed-MLLM. By manually feeding hundreds of identical inputs into MindSpeed-MLLM and the transformer-based forward code, the Mean Absolute Error (MAE) and Mean Relative Error (MRE) are both within five thousandths.

Next, we conducted precision checks through training tasks. In the COCO data set, as illustrated in Figure 5, the loss curves of all frameworks exhibited very similar trends, although with minor discrepancies. These differences are deemed acceptable, as the frameworks utilize distinct data loading modules and workflows.

To further validate consistency, we adapted our custom data loader to MindSpeed-MM, ensuring full alignment of the data load logic, and conducted comparative experiments on the in-house Slow Thinking dataset. As illustrated in Figures 6 and 7, after standardizing the data loading process, the loss values of MindSpeed-MLLM and MindSpeed-MM showed a near-perfect overlap, with the loss difference confined to within a percentile range. This marginal discrepancy is considered acceptable, given the inherent different associated with the fused operators employed, as well as the randomness with computation and communication processes.

In the final benchmark evaluation, under the same training set, the models trained using Llama-Factory and MindSpeed-MLLM respectively achieved comparable results on the general visual and text benchmarks, with an error margin within ± 1.5 .



Figure 6: loss decline trend on the in-house slow thinking dataset.




Figure 7: loss difference between MindSpeed-MLLM and the original MindSpeed-MM.

B DATA CURATION

The MindVL training corpus contains 447 billion diverse and high-quality tokens used for three training stages. The data is categorized according to target capabilities, and the curation process for each category is detailed in the following subsections.

B.1 WARM-UP DATA

The MindVL warm-up corpus contains 256 billion diverse, high-quality tokens, including Generic Image-Text Pairs, Optical Character Recognition (OCR), Visual Grounding & Counting, Science, Technology, Engineering, and Mathematics (STEM), and Graphical User Interface (GUI).

B.1.1 GENERIC IMAGE-TEXT PAIRS

Web-sourced image-text pairs, which include alt texts, captions, and surrounding contextual text, are now available at an unprecedented scale. With billions of examples, they showcase remarkable diversity across visual and textual concepts. However, such data is inherently noisy, often containing texts that are irrelevant or factually inaccurate relative to the corresponding images, and it frequently exhibits class imbalance.

To address these challenges, we have adopted a series of filtering measures, including image-based criteria, text-based criteria, aesthetic scoring, pornography and violence scoring, watermark scoring, CLIP scoring (Radford et al., 2021), and URL filtering. For data with inadequate image-text relevance, we use a model (Bai et al., 2025; Chen et al., 2024) to recapture the images.

Furthermore, to cover more visual concepts and preserve as many types of long-tail visual concepts as possible, we perform clustering on the images within the filtered image-text pairs. We then manually inspect the data in each cluster and map these clusters to the data’s category labels. For clusters that represent overly broad concepts, we conduct re-clustering. Finally, we obtain the results of hierarchical clustering: a larger variance within a cluster indicates greater diversity. We perform data sampling based on variance—clusters with larger variance undergo more extensive sampling, while clusters with smaller variance are sampled less.

Table 11: Warm-up Phase Dataset Composition

Category	Dataset	Total Num	Percentage (%)
Caption	Caption-CN-Rawcaption	132,900,000	21.92
Caption	Caption-EN-Recaption	102,500,000	16.90
Caption	Caption-EN-Rawcaption	42,041,409	6.93
Caption	CapsFusion-en	35,840,000	5.91
Caption	Caption-CN-Recaption	25,300,000	4.17
Caption	HiVision	1,800,000	0.30
Caption Subtotal		338,381,409	55.81
Grounding	UMG-en	42,240,000	6.97
Grounding	Grounding	35,000,000	5.77
Grounding	UMG-zh	14,080,000	2.32
Grounding	GRIT-en	8,448,000	1.39
Grounding Subtotal		99,768,000	16.45
OCR	WuKong-Text-zh	33,280,000	5.49
OCR	LAION-Text-en	14,336,000	2.36
OCR	Document-OCR	17,000,000	2.80
OCR	Scene-OCR	17,000,000	2.80
OCR	IIT-CDIP-en	3,840,000	0.63
OCR	Rendered-Text-en	1,920,000	0.32
OCR	PDF-en	2,560,000	0.42
OCR	PDF-zh	2,560,000	0.42
OCR	arXiv-en	2,304,000	0.38
OCR	DOCX-en	960,000	0.16
OCR	DOCX-zh	960,000	0.16
OCR	README	1,280,000	0.21
OCR	Scene-Text	560,000	0.09
OCR Subtotal		98,000,000	16.16
Table&Chart	Table-Data	10,000,000	1.65
Table&Chart	Chart-SFT-en	1,440,000	0.24
Table&Chart	PubTables-en	120,000	0.02
Table&Chart	PubTables-zh	120,000	0.02
Table&Chart	MMC-Align-en	420,000	0.07
Table&Chart	MMC-Instruct-en	420,000	0.07
Table&Chart Subtotal		12,520,000	2.06
STEM	STEM-Caption-0729-en	2,230,000	0.37
STEM	STEM-Caption-0729-zh	2,230,000	0.37
STEM	STEM-Shuffle	1,000,000	0.16
STEM Subtotal		5,460,000	0.90
GUI	Screen-UI-zh	360,000	0.06
GUI	Screen-UI-en	120,000	0.02
GUI	Screen-QA-zh	120,000	0.02
GUI	Screen-Ref-zh	120,000	0.02
GUI	Screen-Navi-zh	120,000	0.02
GUI Subtotal		840,000	0.14
Total		606,221,409	100.00

B.1.2 OPTICAL CHARACTER RECOGNITION (OCR)

OCR data consists of two types, document OCR data and scene OCR data, each following a structured processing workflow. For document OCR data, the process starts with acquiring and initially screening source PDF data, then processes PDFs page-by-page via a pipeline to generate image-text pairs and image-text interleaved formats. After extracting these data (with image markers removed from pair data and text-only MD files filtered from interleaved data), rule-based cleaning addresses issues like garbled characters or abnormal spaces. For scene OCR data, pre-processed tar packages are first sent to the Focus OCR model (Liu et al., 2024b) to generate OCR results; the same tar packages are then processed by the PaddleOCR model (Cui et al., 2025), with results from both models intersected to form production-line OCR results. In addition to the aforementioned process, we have also calculated the coverage rate of words and characters in OCR text, and conducted targeted supplementation for data with relatively low coverage. In addition, we have also incorporated data such as small-language content, ancient books, and calligraphy works.

B.1.3 TABLE & CHART

For raw tables, we perform layout detection and content recognition separately. Specifically, we use the RapidTable (Team) and LORE (Xing et al., 2023; Long et al., 2025) models for layout detection and layout judgment, and PaddleOCR for content recognition. Subsequently, we integrate the results and conduct post-processing for HTML formatting. Finally, we convert the table images into HTML format. Alongside the previously mentioned data, we also collect open data to enhance the ability to interpret tables and charts. For tables, we use the PubTables-1M dataset (Smock et al., 2022), including both its original English version and a translated Chinese version, to gather table recognition data. For charts, we employ chart-to-table conversion and chart-based QA data from existing datasets, including MMC (Liu et al., 2023) and ChartSFT (Meng et al., 2024).

B.1.4 VISUAL GROUNDING

First, we perform recaptioning on the crawled and open-source images. We then use the Florence-2-large model (Xiao et al., 2024) to conduct caption to phrase grounding processing. After obtaining object bounding boxes and region captions, we employ SAM (Kirillov et al., 2023) to generate point annotations. In addition, we utilize GrIT-20M (Peng et al., 2023), a synthetic caption dataset with additional location labels for major visual elements, to enhance the grounding capability. We incorporate referring and grounding data from UMG-41M (Shi et al., 2024). Specifically, each region is randomly assigned to either a referring task or a grounding task. In the referring task, we provide the model with a bounding box to generate a caption for that specific region, while in the grounding task, we reverse this by using the caption to predict the corresponding bounding box.

B.1.5 SCIENCE, TECHNOLOGY, ENGINEERING, AND MATHEMATICS (STEM)

To enhance the model’s disciplinary knowledge, we incorporated a diverse collection of data across various STEM domains, obtained through both crawling and manual annotation. We referred to the disciplinary classifications of MMMU (Yue et al., 2024) and K12, extracted such data from a vast volume of PDFs, and also collected publicly available data. Additionally, we used a model to re-annotate the sourced disciplinary images. Ultimately, we obtained nearly 5 million data samples.

B.1.6 GRAPHICAL USER INTERFACE (GUI)

For GUI data processing, we first perform pre-annotation via GPT-4V (OpenAI, 2023), which involves tasks such as Referring&Grounding (for icons, text, and coitems), ScreenQA, and Navigation. Next, human verification is conducted by human annotators, focusing on Referring-coitem, QA, and the answers for Navigation. After that, a sequence of data post-processing operations is carried out: first, merging with original annotations; second, using UI-Hawk (Zhang et al., 2024b) and GLM-4V (GLM et al., 2024) to annotate referring-icon, and UI-Hawk to annotate navigation; third, employing GPT-4o (Hurst et al., 2024) to annotate Navigation actions; fourth, merging the results of these three models and performing consistency checks; finally, rewriting the answer style with Qwen2.5-7B-Instruct to ensure consistency and linguistic quality.

The data used during the warm-up phase is shown in the Table 11. Note that the "Total num" in the Table 11 is an approximate value. 606,221,409 samples are about 256B tokens. Apart from Caps-Fusion, UMG, GRIT, PubTables, ChartSFT, and MMC, all other datasets are either self-constructed or have undergone further processing of open-source data. The 16B tokens of data used in MindVL-671B-A37B are uniformly sampled from all warm-up datasets.

B.2 MULTITASK TRAINING DATA

Multitask training data consists of interleaved image-text (Table 12), visual instruction (Table 14, about 80B tokens), web2code and text instruction (Table 13), totally 179B tokens.

B.2.1 INTERLEAVED IMAGE-TEXT DATA

For interleaved image-text data, we primarily crawl Chinese websites, English websites, news websites, and Wikipedia from the internet. We then use model-based methods to filter the crawled data. First, we apply CLIP scores to filter the data according to image-text relevance. Next, we task the model with scoring the image-text interleaved data across multiple dimensions, including question-and-answer quality, relevance between images and questions, complementarity between answers, images and questions, and balance of information density. The model is also required to provide justifications for these scores. Finally, we use these scores to filter the image-text interleaved data.

Table 12: Data Statistics and Proportion of Interleaved Image-Text Data. "GPT4o" / "QWEN" denotes we use GPT4o / Qwen2.5-VL 72B filter the crawled data.

Data Category	Single-image Count	Multi-image Count	Total Count	Proportion (%)
HTMLENGPT4o	2994750	2994750	5989500	12.24
HTMLENQWEN	356942	356942	713884	1.46
HTMLZHGPT4o	763016	763016	1526032	3.12
HTMLZHQWEN	8096192	8096192	16192384	33.09
OBELICSGPT4o	3339440	3339440	6678880	13.65
News Webpages	4167856	4167856	8335712	17.03
Knowledge	—	—	7500000	15.33
Wikiweb2m	—	—	2000000	4.09
Total	—	—	48936392	100.00

For other types of multitask training data, we have collected a wide range of open source data, as presented in Table 14. Within this dataset, we used Qwen2.5VL-72B to reannotate certain portions, such as those whose names begin with "Qwen". In addition, we also sampled some data from the warm-up phase and incorporated it into the training process at this stage. 200M web2code data is also used in this stage. Furthermore, we utilized language data produced by DeepSeek R1 to preserve the MLLM's linguistic abilities as shown in Table 13, with the ratio of multimodal to language data being approximately 8:2 at this stage.

Table 13: Data Statistics and Proportion of text instruction.

Data Source	Number	Proportion (%)
Competition Mathematics	345606	16.06
Code	113064	5.25
Math	296500	13.78
General QA	1316361	61.18
Tulu Instruction	19118	0.89
Multi-turn Conversation	61114	2.84
Total	2151763	100.00

Table 14: Data Statistics and Proportion of Visual Instruction

No.	Data Set	Number	Percentage (%)
1	Cauldron/ai2d	2434	0.0064
2	Cauldron/cocoqa	46287	0.1214
3	Cauldron/datikz	47974	0.1258
4	Cauldron/diagram-image-to-text	300	0.0008
5	Cauldron/docvqa	10189	0.0267
6	Cauldron/dvqa	199937	0.5242
7	Cauldron/figureqa	100000	0.2622
8	Cauldron/finqa	5276	0.0138
9	Cauldron/geomverse	9303	0.0244
10	Cauldron/hateful-memes	8500	0.0223
11	Cauldron/hitab	2500	0.0066
12	Cauldron/iam	5663	0.0148
13	Cauldron/iconqa	27307	0.0716
14	Cauldron/infographic-vqa	2118	0.0056
15	Cauldron/intergps	1280	0.0034
16	Cauldron/localized-narratives	199998	0.5244
17	Cauldron/mapqa	37417	0.0981
18	Cauldron/mimic-cgd	70939	0.1860
19	Cauldron/multihiertr	7619	0.0200
20	Cauldron/nlvr2	50426	0.1322
21	Cauldron/plotqa	157007	0.4117
22	Cauldron/raven	42000	0.1091
23	Cauldron/rendered-text	10000	0.0262
24	Cauldron/robut-sqa	8514	0.0223
25	Cauldron/robut-wikisql	74989	0.1966
26	Cauldron/robut-wtq	38246	0.1003
27	Cauldron/scienceqa	4976	0.0130
28	Cauldron/screen2words	15730	0.0412
29	Cauldron/spot-the-diff	8566	0.0225
30	Cauldron/tabmwp	22722	0.0596
31	Cauldron/tallyqa	98680	0.2587
32	Cauldron/tat-qa	2199	0.0058
33	Cauldron/textcaps	21953	0.0576
34	Cauldron/textvqa	21953	0.0576
35	Cauldron/tqa	1493	0.0039
36	Cauldron/vistext	9969	0.0261
37	Cauldron/visual7w	14366	0.0377
38	Cauldron/visualmrc	3027	0.0079
39	Cauldron/vqarad	313	0.0008
40	Cauldron/vsr	2157	0.0056
41	Cauldron/websight	10000	0.0262
42	ChartQA/augmented	15474	0.0406
43	ChartQA/cap	18207	0.0477
44	ChartQA/cot	18215	0.0478
45	ChartQA/gemini-v	16393	0.0430
46	ChartQA/gemini-v-cot	16393	0.0430
47	ChartQA/human	3699	0.0097
48	ChartSFT/unichart	120554	0.3161
49	CoSyn-400K/chart	116807	0.3063
50	CoSyn-400K/chemical	8935	0.0234
51	CoSyn-400K/circuit	10463	0.0274
52	CoSyn-400K/diagram	34956	0.0916
53	CoSyn-400K/document	71275	0.1869
54	CoSyn-400K/graphic	26961	0.0707
55	CoSyn-400K/math	66707	0.1749

Data Statistics and Proportion of Visual Instruction (Continue)

No.	Data Set	Number	Percentage (%)
56	CoSyn-400K/music	11962	0.0314
57	CoSyn-400K/nutrition	6924	0.0181
58	CoSyn-400K/table	46511	0.1220
59	CoSyn-point	19385	0.0508
60	DocVQA/cap	9892	0.0259
61	DocVQA/cot	11938	0.0313
62	DocVQA/gemini-v	11182	0.0293
63	DocVQA/gemini-v-cot	11182	0.0293
64	DocVQA/human	10194	0.0267
65	InfoVQA/cap	4290	0.0112
66	InfoVQA/cot	6060	0.0159
67	InfoVQA/gemini-v	5187	0.0136
68	InfoVQA/gemini-v-cot	5187	0.0136
69	InfoVQA/human	4406	0.0116
70	M4-Instruct/3D-LLM-3-datasets	49890	0.1308
71	M4-Instruct/AESOP	6915	0.0181
72	M4-Instruct/ALFRED	22565	0.0592
73	M4-Instruct/Birds-to-Words	14281	0.0374
74	M4-Instruct/CLEVR-Change	3885	0.0102
75	M4-Instruct/FlintstonesSV	22341	0.0586
76	M4-Instruct/HQ-Edit	50000	0.1311
77	M4-Instruct/HQ-Edit-Diff	7000	0.0184
78	M4-Instruct/IEEdit	3456	0.0091
79	M4-Instruct/MIT-States-PropertyCoherence	1900	0.0050
80	M4-Instruct/MIT-States-StateCoherence	1900	0.0050
81	M4-Instruct/MagicBrush	14249	0.0374
82	M4-Instruct/MagicBrush-Diff	6698	0.0176
83	M4-Instruct/PororoSV	12299	0.0322
84	M4-Instruct/RAVEN	35000	0.0918
85	M4-Instruct/RecipeQA-ImageCoherence	8699	0.0228
86	M4-Instruct/ScanQA	25563	0.0670
87	M4-Instruct/TQA	8249	0.0216
88	M4-Instruct/VISION	9900	0.0260
89	M4-Instruct/VIST	26026	0.0682
90	M4-Instruct/WebQA	9338	0.0245
91	M4-Instruct/coinstruct	50000	0.1311
92	M4-Instruct/contrastive-caption	25240	0.0662
93	M4-Instruct/dreamsim	15941	0.0418
94	M4-Instruct/iconqa	34603	0.0907
95	M4-Instruct/imagecode	16594	0.0435
96	M4-Instruct/multi-vqa	4993	0.0131
97	M4-Instruct/nextqa	3870	0.0102
98	M4-Instruct/nlvr2	86373	0.2265
99	M4-Instruct/star	3032	0.0079
100	M4-Instruct/twitter-post	5734	0.0150
101	Markdown/arxiv-v2	200000	0.5244
102	Markdown/docx-en	200000	0.5244
103	Markdown/docx-zh	200000	0.5244
104	Markdown/pubtables-c	200000	0.5244
105	Markdown/pubtables-e	200000	0.5244
106	Markdown/pubtables-o	200000	0.5244
107	Markdown/readme-v2	200000	0.5244
108	Math-PUMA/Synthesis-train-answer	120239	0.3153
109	Math-PUMA/Synthesis-train-instruction	120239	0.3153
110	Math-PUMA/VarsityTutors-pure-text	232959	0.6108

Data Statistics and Proportion of Visual Instruction (Continue)

No.	Data Set	Number	Percentage (%)
111	Math-PUMA/VarsityTutors-with-image	78935	0.2070
112	Math-PUMA/VisualWebInstruct-train-answer	263599	0.6911
113	OCR/iit-cdip-v2-en	200000	0.5244
114	OCR/laion-text-en	50000	0.1311
115	OCR/laion-text-v2-en	200000	0.5244
116	OCR/pdf-en	200000	0.5244
117	OCR/pdf-zh	200000	0.5244
118	OCR/wukong-text-v2-zh	200000	0.5244
119	OCR/wukong-text-zh	50000	0.1311
120	OneVision/ai2d-gpt4v	4833	0.0127
121	OneVision/ai2d-internvl	12372	0.0324
122	OneVision/clevr-math-mathv360k	5249	0.0138
123	OneVision/figureqa-mathv360k	17556	0.0460
124	OneVision/geo3k	2060	0.0054
125	OneVision/geometry3k-mathv360k	9693	0.0254
126	OneVision/geoqa-plus-mathv360k	17131	0.0449
127	OneVision/geos-mathv360k	467	0.0012
128	OneVision/iconqa-mathv360k	22558	0.0591
129	OneVision/iiit5k	1959	0.0051
130	OneVision/image-textualization-filtered	99542	0.2610
131	OneVision/infographic-gpt4v	1951	0.0051
132	OneVision/k12-printing	256605	0.6728
133	OneVision/lrv-chart	1745	0.0046
134	OneVision/lrv-normal-filtered	10456	0.0274
135	OneVision/mapqa-mathv360k	5194	0.0136
136	OneVision/orand-car-a	1968	0.0052
137	OneVision/pmc-vqa-mathv360k	35917	0.0942
138	OneVision/scienceqa-nona-context	19177	0.0503
139	OneVision/super-clevr-mathv360k	8611	0.0226
140	OneVision/tabmwp-mathv360k	22421	0.0588
141	OneVision/unigeo-mathv360k	11918	0.0312
142	OneVision/vision-flan-filtered	186029	0.4877
143	OneVision/vizwiz-mathv360k	6573	0.0172
144	PixMo/ask-model-anything	161737	0.4241
145	PixMo/ask-model-anything-zh	140343	0.3679
146	PixMo/cap	717042	1.8800
147	PixMo/cap-qa	271714	0.7124
148	PixMo/cap-qa-zh	194097	0.5090
149	PixMo/cap-zh	508134	1.3323
150	PixMo/count	36916	0.0968
151	PixMo/docs-charts	116814	0.3063
152	PixMo/docs-diagrams	16551	0.0434
153	PixMo/docs-other	71282	0.1869
154	PixMo/docs-tables	46518	0.1220
155	PixMo/point-explanations	79551	0.2086
156	PixMo/points	2376222	6.2301
157	QwenCaps/PixMo/cap	712401	1.8680
158	QwenCaps/PixMo/cap-zh	642166	1.6837
159	QwenCaps/ai2d	9778	0.0256
160	QwenCaps/allava-laion	480858	1.2608
161	QwenCaps/chart2text	16605	0.0435
162	QwenCaps/chart-to-text	44096	0.1156
163	QwenCaps/chartqa	18317	0.0480
164	QwenCaps/coco	163119	0.4277
165	QwenCaps/docvqa	11166	0.0293

Data Statistics and Proportion of Visual Instruction (Continue)

No.	Data Set	Number	Percentage (%)
166	QwenCaps/flickr30k	31772	0.0833
167	QwenCaps/funsd	149	0.0004
168	QwenCaps/infvqa	4890	0.0128
169	QwenCaps/laion-gpt4v	10962	0.0287
170	QwenCaps/llava	557086	1.4606
171	QwenCaps/llavar-s1	47231	0.1238
172	QwenCaps/ocr-vqa	208190	0.5458
173	QwenCaps/opencqa	7724	0.0202
174	QwenCaps/poie	2999	0.0079
175	QwenCaps/sam	33533	0.0879
176	QwenCaps/scienceqa	4985	0.0131
177	QwenCaps/screen2words	15669	0.0411
178	QwenCaps/sroie	608	0.0016
179	QwenCaps/textvqa	24990	0.0655
180	QwenCaps/vg	105345	0.2762
181	QwenCaps/vision-flan	181189	0.4751
182	QwenCaps/vistext	8816	0.0231
183	QwenCaps/xfund	1043	0.0027
184	QwenQA/Cauldron/cocoqa	27022	0.0708
185	QwenQA/Cauldron/diagram-image-to-text	297	0.0008
186	QwenQA/Cauldron/hateful-memes	3729	0.0098
187	QwenQA/Cauldron/mapqa	22950	0.0602
188	QwenQA/Cauldron/tabmwp	22410	0.0588
189	QwenQA/Cauldron/tallyqa	83671	0.2194
190	QwenQA/Cauldron/textvqa	16059	0.0421
191	QwenQA/Cauldron/tqa	861	0.0023
192	QwenQA/Cauldron/visual7w	13006	0.0341
193	QwenQA/Cauldron/vsr	1650	0.0043
194	QwenQA/ChartQA/augmented	12469	0.0327
195	QwenQA/ChartQA/human	3290	0.0086
196	QwenQA/DocVQA/human	8924	0.0234
197	QwenQA/InfoVQA/human	4092	0.0107
198	QwenQA/PixMo/cap-qa	247256	0.6483
199	QwenQA/PixMo/cap-qa-zh	218868	0.5738
200	QwenQA/ai2d	10593	0.0278
201	QwenQA/allava-laion	484257	1.2697
202	QwenQA/allava-vflan	179020	0.4694
203	QwenQA/aokvqa	39358	0.1032
204	QwenQA/gqa	69198	0.1814
205	QwenQA/okvqa	3442	0.0090
206	QwenQA/raw/ChartQA/augmented	15474	0.0406
207	QwenQA/raw/ChartQA/human	3699	0.0097
208	QwenQA/raw/DocVQA/human	10194	0.0267
209	QwenQA/raw/InfoVQA/human	4394	0.0115
210	QwenQA/raw/ai2d	12413	0.0325
211	QwenQA/raw/aokvqa	76762	0.2013
212	QwenQA/raw/gqa	72140	0.1891
213	QwenQA/raw/okvqa	8996	0.0236
214	QwenQA/raw/scienceqa	12726	0.0334
215	QwenQA/raw/stvqa	18921	0.0496
216	QwenQA/raw/tabfact	13127	0.0344
217	QwenQA/raw/vqav2	82783	0.2171
218	QwenQA/raw/wtq	1679	0.0044
219	QwenQA/scienceqa	11071	0.0290
220	QwenQA/stvqa	14381	0.0377

Data Statistics and Proportion of Visual Instruction (Continue)

No.	Data Set	Number	Percentage (%)
221	QwenQA/tabfact	12648	0.0332
222	QwenQA/vqav2	78718	0.2064
223	QwenQA/wtq	1588	0.0042
224	R1-Vision/PixMo/cap-qa	260144	0.6821
225	R1-Vision/PixMo/cap-qa-zh	218962	0.5741
226	Screen/grd	212472	0.5571
227	Screen/grd-small	212472	0.5571
228	Screen/nav	286364	0.7508
229	Screen/nav-small	286364	0.7508
230	Screen/ref	402742	1.0560
231	Screen/ref-small	402742	1.0560
232	Screen/sqa	417579	1.0948
233	Screen/sqa-small	417579	1.0948
234	Screen/ui-en	50000	0.1311
235	Screen/ui-en-small	50000	0.1311
236	Screen/ui-zh	50000	0.1311
237	Screen/ui-zh-small	50000	0.1311
238	XFUND/zh	149	0.0004
239	XFUND/zh-qa	671	0.0018
240	Allava-caption-laion	505588	1.3256
241	Allava-caption-vflan	182864	0.4794
242	Allava-instruct-laion	505588	1.3256
243	Allava-instruct-vflan	183839	0.4820
244	Aokvqa	76762	0.2013
245	Chartsft	808768	2.1205
246	Cipc-docvqa	5243	0.0137
247	Coavqa-bench	36019	0.0944
248	Coavqa-chat	18004	0.0472
249	Coco-rec	130723	0.3427
250	Coco-reg	130723	0.3427
251	Crohme	9821	0.0257
252	Datikz-v2	94532	0.2479
253	Dcc	97306	0.2551
254	Docmatix	1249967	3.2773
255	E2e-ocr	259511	0.6804
256	E2e-ocr-0324	359996	0.9440
257	Eikie-qa	183639	0.4815
258	Estvqa	17047	0.0447
259	Estvqa-ca	17047	0.0447
260	Estvqa-grounding	17047	0.0447
261	Flickr	148915	0.3904
262	Funsd	149	0.0004
263	Funsd-qa	431	0.0011
264	Geo170k-alignment	60252	0.1580
265	Geo170k-qa-tuning	117205	0.3073
266	Geo-synth-caption-1108	320000	0.8390
267	Geo-synth-qa-1108	320004	0.8390
268	Gqa	72140	0.1891
269	Hme	74502	0.1953
270	Kvqa	24602	0.0645
271	Laion-gpt4v	12356	0.0324
272	Llava-s2	157712	0.4135
273	Llarar-s2	19800	0.0519
274	Lrv-instruction	180722	0.4738
275	Lvis-instruct4v	222711	0.5840

Data Statistics and Proportion of Visual Instruction (Continue)

No.	Data Set	Number	Percentage (%)
276	M2e	99956	0.2621
277	Mavis-900k	458042	1.2009
278	Mavis-caps	599748	1.5725
279	Mlhme	30000	0.0786
280	Mmc-non-arxiv-2	21103	0.0553
281	Mmc-non-arxiv-3	88917	0.2331
282	Noahcaps-en	14956	0.0392
283	Noahcaps-zh	44394	0.1164
284	Objects365	1737995	4.5568
285	Ocrvqa	207572	0.5442
286	Okvqa	8996	0.0236
287	Petal-wiki	2396596	6.2836
288	Poie	2250	0.0059
289	Pointqa-general	82254	0.2157
290	Pointqa-local	10403	0.0273
291	Pointqa-looktwice	22838	0.0599
292	Sharegpt4o-image	57289	0.1502
293	Sharegpt4o-video	2111	0.0055
294	Sharegpt4v	102025	0.2675
295	Shikra	5814	0.0152
296	Slidecaps	34810	0.0913
297	Slidevqa	10341	0.0271
298	Sroie	500	0.0013
299	Stvqa	18921	0.0496
300	Svit	136349	0.3575
301	Tabfact	13182	0.0346
302	Tech4all	8256	0.0216
303	Textocr-gpt4v	25114	0.0658
304	Tiku/math-4o-zh	386025	1.0121
305	Tiku/multidisciplinary-en	2220078	5.8207
306	Unimer	987289	2.5885
307	VCR	95036	0.2492
308	Vflan-4o-zh	35899	0.0941
309	VG-rem	611465	1.6032
310	Viquae	1190	0.0031
311	VQA-cdip-anneal	753921	1.9767
312	VQAv2	82783	0.2171
313	WTQ	1679	0.0044
Total		38140772	100.0000

B.3 SUPERVISED FINE-TUNING DATA

During the initial phase of SFT data construction, our objective is to equip the model to handle a broad spectrum of application scenarios. To this end, we develop a model capability taxonomy based on the classification of traditional visual tasks and practical application requirements of vision-language models. Guided by this taxonomy, we sample from open-source multi-task training data to supplement high-quality instruction data. We use a pretrained model to classify the open-source data and extract the required categories. For data where image and question quality are high but answer quality is poor, we re-annotate the answers using a model and perform manual verification. Additionally, based on category guidelines, we generate targeted instruction data for certain categories to further enhance the model’s capabilities. Furthermore, we utilized language data produced by DeepSeek R1 to preserve the MLLM’s linguistic abilities, with the ratio of multimodal to language data being approximately 1:1 at this stage.

Table 15: Detailed Statistics of SFT Data

No.	Dataset Name	Number	Percentage (%)
1	Competition Mathematics	150000	2.4269
2	Code	50000	0.8089
3	Math	200000	3.2359
4	General QA	300000	4.8539
5	Tulu Instruction	19118	0.3093
6	Multi-turn Conversation	61114	0.9888
7	MegaScience	1250000	20.2244
8	OpenHermes2-5	200000	3.2359
9	QwenCaps/PixMo/cap	712401	11.5263
10	QwenCaps/PixMo/cap-zh	10000	0.1618
11	QwenCaps/Allava-Laion	10000	0.1618
12	QwenCaps/COCO	10000	0.1618
13	QwenCaps/Flickr30k	10000	0.1618
14	QwenCaps/Laion-GPT4V	10000	0.1618
15	QwenCaps/LLaVA	10000	0.1618
16	QwenCaps/LLaVaR-S1	10000	0.1618
17	QwenCaps/OCR-VQA	10000	0.1618
18	QwenCaps/SAM	10000	0.1618
19	QwenCaps/VG	10000	0.1618
20	QwenCaps/Vision-FLAN	10000	0.1618
21	DCC	10000	0.1618
22	NoahCaps-EN	10000	0.1618
23	NoahCaps-ZH	10000	0.1618
24	ShareGPT4O-Image	10000	0.1618
25	ShareGPT4O-Video	10000	0.1618
26	Tech4All	10000	0.1618
27	OneVision/CLEVR-Math-MathV360K	10000	0.1618
28	OneVision/Super-CLEVR-MathV360K	10000	0.1618
29	PixMo/Ask-Model-Anything	161737	2.6168
30	PixMo/Ask-Model-Anything-ZH	10000	0.1618
31	PixMo/Count	10000	0.1618
32	QwenQA/Cauldron/COCOQA	27022	0.4372
33	QwenQA/Cauldron/TallyQA	83671	1.3538
34	QwenQA/Cauldron/Visual7W	13006	0.2104
35	QwenQA/Cauldron/VSR	1650	0.0267
36	QwenQA/PixMo/Cap-QA	247256	3.9990
37	QwenQA/PixMo/Cap-QA-ZH	10000	0.1618
38	QwenQA/Allava-Laion	484257	7.8350
39	QwenQA/Allava-VFLAN	179020	2.8965
40	QwenQA/AOKVQA	39358	0.6368
41	QwenQA/GQA	69198	1.1196
42	QwenQA/OKVQA	3442	0.0557
43	QwenQA/VQAv2	78718	1.2736
44	QwenQA/Raw/AOKVQA	10000	0.1618
45	QwenQA/Raw/GQA	10000	0.1618
46	QwenQA/Raw/OKVQA	10000	0.1618
47	QwenQA/Raw/VQAv2	10000	0.1618
48	VFLAN-4O-ZH	10000	0.1618
49	Cauldron/IconQA	10000	0.1618
50	Cauldron/Mimic-CGD	10000	0.1618
51	Cauldron/NLVR2	10000	0.1618
52	Cauldron/Raven	10000	0.1618
53	Cauldron/Spot-the-Diff	10000	0.1618
54	PixMo/Points	10000	0.1618
55	COCO-Rec	130723	2.1150

Detailed Statistics of SFT Data (Continued).

No.	Dataset Name	Number	Percentage (%)
56	COCO-Reg	130723	2.1150
57	Flickr	10000	0.1618
58	Objects365	10000	0.1618
59	Cauldron/IAM	10000	0.1618
60	CoSyn-400K/Chart	10000	0.1618
61	CoSyn-400K/Chemical	10000	0.1618
62	CoSyn-400K/Circuit	10000	0.1618
63	CoSyn-400K/Diagram	10000	0.1618
64	CoSyn-400K/Document	10000	0.1618
65	CoSyn-400K/Graphic	10000	0.1618
66	CoSyn-400K/Music	10000	0.1618
67	CoSyn-400K/Nutrition	10000	0.1618
68	CoSyn-400K/Table	10000	0.1618
69	OneVision/IIIT5K	10000	0.1618
70	OneVision/K12-Printing	10000	0.1618
71	OneVision/LRV-Chart	10000	0.1618
72	OneVision/Orand-Car-A	10000	0.1618
73	QwenCaps/ChartQA	18317	0.2964
74	QwenCaps/Chart2Text	16605	0.2687
75	QwenCaps/Chart-to-Text	44096	0.7135
76	QwenCaps/DocVQA	11166	0.1807
77	QwenCaps/FUNSD	149	0.0024
78	QwenCaps/InfoVQA	4890	0.0791
79	QwenCaps/OpenCQA	7724	0.1250
80	QwenCaps/POIE	2999	0.0485
81	QwenCaps/SROIE	608	0.0098
82	QwenCaps/TextVQA	24990	0.4043
83	QwenCaps/Vistext	8816	0.1426
84	QwenCaps/XFUND	1043	0.0169
85	QwenQA/Cauldron/Diagram-Image-to-Text	297	0.0048
86	QwenQA/Cauldron/Hateful-Memes	3729	0.0603
87	QwenQA/Cauldron/TabMWP	22410	0.3626
88	QwenQA/Cauldron/TextVQA	16059	0.2598
89	QwenQA/ChartQA/Human	3290	0.0532
90	QwenQA/ChartQA/Augmented	12469	0.2017
91	QwenQA/DocVQA/Human	8924	0.1444
92	QwenQA/InfoVQA/Human	4092	0.0662
93	QwenQA/STVQA	14381	0.2327
94	QwenQA/TabFact	12648	0.2046
95	QwenQA/WTQ	1588	0.0257
96	QwenQA/Raw/STVQA	10000	0.1618
97	QwenQA/Raw/TabFact	10000	0.1618
98	QwenQA/Raw/WTQ	10000	0.1618
99	XFUND/ZH-QA	10000	0.1618
100	CIPC-DocVQA	10000	0.1618
101	CoAVQA-Chat	10000	0.1618
102	CoAVQA-Bench	10000	0.1618
103	E2E-OCR-0324	359996	5.8245
104	ESTVQA	10000	0.1618
105	ESTVQA-Grounding	10000	0.1618
106	FUNSD-QA	10000	0.1618
107	POIE	10000	0.1618
108	SlideCaps	10000	0.1618
109	SlideVQA	10000	0.1618
110	SROIE	10000	0.1618

Detailed Statistics of SFT Data (Continued).

No.	Dataset Name	Number	Percentage (%)
111	VQA-CDIP-SFT	10000	0.1618
112	QwenCaps/Screen2Words	10000	0.1618
113	Cauldron/VisualMRC	10000	0.1618
114	Cauldron/WebSight	10000	0.1618
115	CoSyn-Point	10000	0.1618
116	Screen/Nav	10000	0.1618
117	Screen/Ref	10000	0.1618
118	Screen/SQA	10000	0.1618
119	Cauldron/GeomVerse	10000	0.1618
120	CoSyn-400K/Math	66707	1.0793
121	OneVision/Geometry3K-MathV360K	10000	0.1618
122	OneVision/GeoQA-Plus-MathV360K	10000	0.1618
123	OneVision/Geos-MathV360K	10000	0.1618
124	OneVision/IconQA-MathV360K	10000	0.1618
125	OneVision/PMC-VQA-MathV360K	10000	0.1618
126	OneVision/TabMWP-MathV360K	10000	0.1618
127	OneVision/Unigeo-MathV360K	10000	0.1618
128	OneVision/Geo3K	10000	0.1618
129	OneVision/Mavis-Math-Metagen	10000	0.1618
130	OneVision/Mavis-Math-Rule-Geo	10000	0.1618
131	CROHME	10000	0.1618
132	Geo170K-Alignment	10000	0.1618
133	Geo170K-QA-Tuning	10000	0.1618
134	HME	10000	0.1618
135	M2E	10000	0.1618
136	MLHME	10000	0.1618
137	Unimer	10000	0.1618
138	QwenCaps/AI2D	9778	0.1582
139	QwenCaps/ScienceQA	4985	0.0807
140	QwenQA/Cauldron/MapQA	22950	0.3713
141	QwenQA/Cauldron/TQA	861	0.0139
142	QwenQA/AI2D	10593	0.1714
143	QwenQA/ScienceQA	11071	0.1791
Total		6180645	100.0000

C EXPERIMENTS

C.1 RESULTS OF TEST-TIME RESOLUTION SEARCH

In this Section, we present detail evaluation results with test-time image resolution search strategy. On MME dataset, as shown in Figure 8, models trained with different resolution lengths yield consistent evaluation results when min-pixels is less than 16. In addition, down-sampling images on MME results in a significant improvement in model performance. On OCRBench dataset, as shown in Figure 9, when min-pixels is set to 16 or 32, the model performs better than when it is set to 4. The reason is that too small input images are out of distribution in OCR task for our model. On the other hand, when min-pixels is set to 64, the model performance degrades. The reason is that the image blurring effect caused by up-sampling images has a negative impact on the model’s visual perception. On DocVQA and InfoVQA datasets, as shown in Figure 10 and Figure 11, min-pixels variable has no detectable impact on the evaluation results since most of the images in these two datasets have a resolution higher than min-pixels. Furthermore, the optimal inference max-pixels on these two datasets falls within the range of 2560–3072.

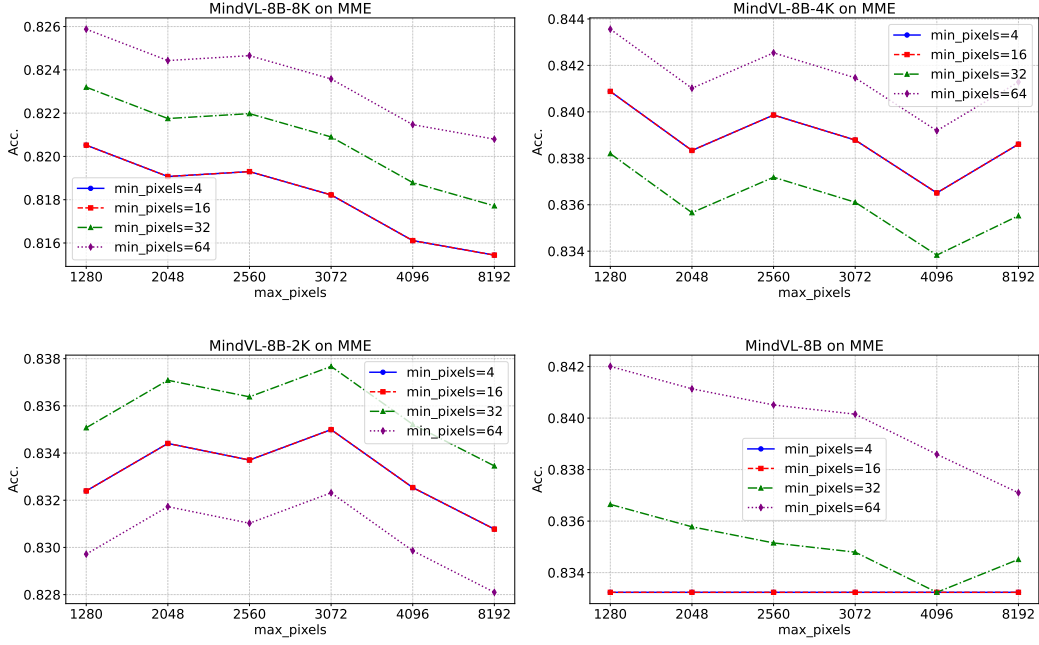


Figure 8: Evaluation results on MME of MindVL with different input image resolutions.

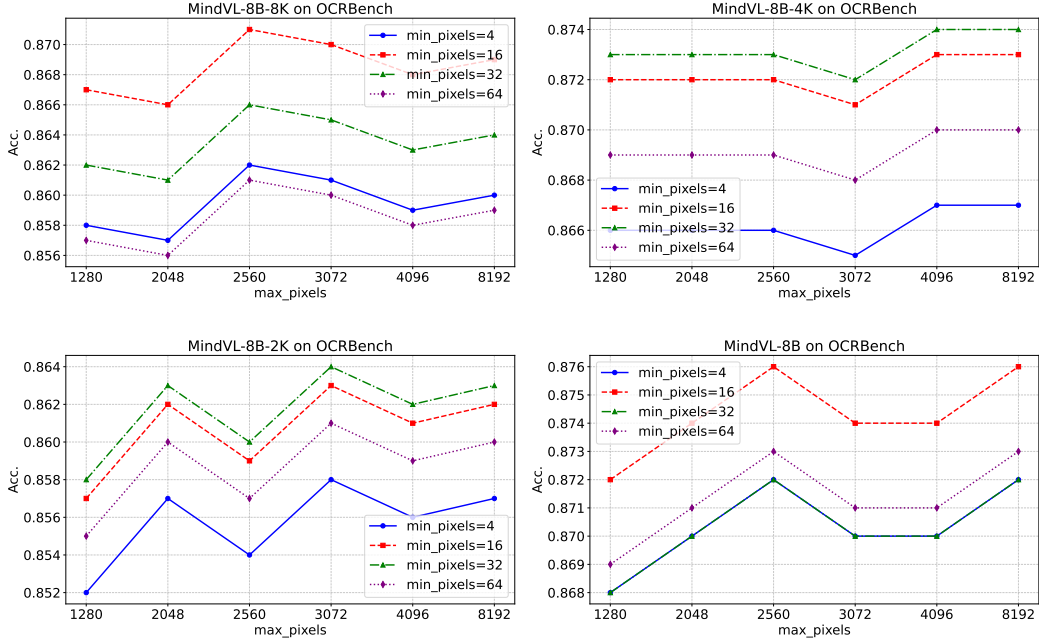


Figure 9: Evaluation results on OCRBench of MindVL with different input image resolutions.

C.2 INTEGRATION OF MULTITASK DATA

We also present the experimental results of the model trained with different data recipes. As shown in Table 16, MindVL-8B-V1Data model is trained with introduced data of Video, Captions, OCR, GUI, Math, Knowledge, Grounding and Genreal VQA introduced in Section 4.2. For MindVL-8B-V2Data model, Internleaved image-text data are added for training. For MindVL-8B-V3Data model, STEM data and additional Captions, OCR, GUI, Grounding and Genreal VQA data are added




Figure 10: Evaluation results on DocVQA of MindVL with different input image resolutions.




Figure 11: Evaluation results on InfoVQA of MindVL with different input image resolutions.

for training. By comparing the results of MindVL-8B-V1Data and MindVL-8B-V2Data, we can observe that adding Internleaved image-text data for training can significantly improve performance over all benchmarks. It is because that interleaved image-text data not only contains complex logical relationships between text and images, but also covers a wide range of domain knowledge—both of which are highly beneficial for enhancing the multimodal model’s cognitive capabilities. Moreover, by adding more refined and diverse Captions, OCR, GUI, Grounding and Genreal VQA data to train

the model MindVL-8B-V3Data, the performance on multiple benchmark datasets has been further improved, such as MME, MMBench and InfoVQA.

Table 16: Ablation study results of multitask data.

Model	MME	MMBench	OCRBench	DocVQA	ChartQA	InfoVQA	Overall
MindVL-8B-V1Data	79.8	81.5	83.3	92.1	83.7	73.2	82.3
MindVL-8B-V2Data	82.4	83.4	88.2	94.7	87.2	79.9	86.0
MindVL-8B-V3Data	84.1	84.3	87.6	94.7	87.2	81.1	86.5

C.3 ABLATION STUDY ON TRAINING PHASES

Table 17: Ablation study results on training phases.

Model	MME	MMBench	OCRBench	DocVQA	ChartQA	InfoVQA	Overall
MindVL-8B-4Phase	80.8	82.3	83.6	92.7	86.4	76.9	83.8
MindVL-8B-2Phase*	79.4	81.8	84.9	93.5	85.2	75.8	83.4
MindVL-8B-2Phase	81.6	82.4	87.0	94.2	86.8	79.6	85.3
MindVL-8B-3Phase	84.1	84.3	87.6	94.7	87.2	81.1	86.5


As mentioned in Section 5.2, MindVL-8B undergoes a three-phase training pipeline: warm-up, multi-task learning, and SFT. We denote the model trained on such recipe as MindVL-8B-3Phase. Table 17 shows the results of ablation study on training phases. MindVL-8B-4Phase model is trained on a four-phase training pipeline: warm-up, pretrain, multi-task learning, and SFT. In the pretrain phase, both ViT and MLP adaptor are trained. Our ultimate training recipe has only three-phase: warm-up, multi-task learning and SFT, achieves best performance by merging training data in warm-up and pretrain phases of MindVL-8B-4Phase. By changing the four-phase training recipe to a three-phase training recipe, we have not only improved training efficiency but also enhanced the model performances. Such results indicate that training a MLP adaptor can also align ViT and LLMs effectively.

We also conducted experiments with two-phase training recipes, MindVL-8B-2Phase* undergoes warm-up and SFT, and MindVL-8B-2Phase undergoes warm-up and multi-task learning with the training data in corresponding phase. The results show that MindVL-8B-2Phase outperformance MindVL-8B-2Phase*, which indicate that high-quality multi-task data is crucial for enhancing the capabilities of multimodal models.

D EXAMPLES

We list some samples whose answers are predicted by our MindVL-671B-37B.

Image:



Question:

从每个盒子里取出卡片，这样最后每个盒子里只有一张卡片，不同的盒子里有不同字母的卡片。问：第1盒中还剩下哪张牌？

Response:

初始理解题目

我们有五个盒子，每个盒子里面有一些卡片，卡片上标有不同的字母。初始状态如下：

- 盒子1：E, S, X, T

- 盒子2：E, T, I, S

- 盒子3：T, S, E

- 盒子4：E

- 盒子5：E, S

****目标**：**从每个盒子中取出一些卡片，使得最终每个盒子中****只剩下一张卡片****，并且****所有盒子中剩下的卡片上的字母都是不同的****。

分析当前状态

首先，我们需要注意到盒子4已经满足条件，因为它只有一张卡片E。因此，盒子4的E必须保留，不能被取出。这意味着其他盒子中不能保留E，因为每个字母只能出现一次。

确定保留的字母

我们需要为每个盒子选择一个字母，确保：

1. 每个盒子一个字母。

2. 五个字母互不相同。

3. 盒子4的E已经被保留，其他盒子不能保留E。

逐步选择字母

****盒子4**：**已经确定为E。

****盒子5**：**有E和S。E已经被盒子4保留，所以只能选择S。因此，盒子5保留S。

现在已保留的字母：E（盒子4），S（盒子5）。

****盒子3**：**有T, S, E。S和E已经被保留，所以只能选择T。因此，盒子3保留T。

已保留的字母：E, S, T。

****盒子2**：**有E, T, I, S。E, T, S已经被保留，所以只能选择I。因此，盒子2保留I。

已保留的字母：E, S, T, I。

****盒子1**：**有E, S, X, T。E, S, T已经被保留，所以只能选择X。因此，盒子1保留X。

已保留的字母：E, S, T, I, X。

验证

让我们验证一下每个盒子保留的字母是否满足条件：

- 盒子1：X

- 盒子2：I

- 盒子3：T

- 盒子4：E

- 盒子5：S

所有字母X, I, T, E, S都是不同的，且每个盒子只有一张卡片。

可能的其他选择

我们需要确认是否存在其他可能的保留方式。例如：

如果盒子5保留E，那么盒子4必须保留其他字母，但盒子4只有E，所以无法改变。因此，盒子5必须保留S。

类似地，盒子3必须保留T，因为S和E已经被保留。

盒子2必须保留I，因为E, T, S已经被保留。

盒子1必须保留X，因为E, S, T已经被保留。



因此，这种保留方式是唯一的。

结论

根据以上分析，****第1盒中还剩下的牌是X****。

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

Image:



Question:

这会儿我买50块钱的菜能省多少钱？

Response:

根据图片中的时段表，当前时间为19:56，属于19:31至20:00的“全场七折”时段。因此，购买50元的菜可以享受七折优惠。

计算节省的金额：
原价：50元
折扣后价格：50元 × 0.7 = 35元
节省金额：50元 - 35元 = 15元

所以，这会儿买50块钱的菜能省15元。

Image:



Question:

帮我解下这道题

Response:

要得到函数 $y = 3 \cos(2x - \frac{\pi}{3})$ 的图象，只需将函数 $y = 3 \cos 2x$ 的图象进行平移。我们需要确定平移的方向和距离。

解题步骤：

1. **理解函数变换**：

- 函数 $y = 3 \cos(2x - \frac{\pi}{3})$ 可以看作是 $y = 3 \cos 2x$ 的相位移动。
- 一般形式为 $y = A \cos(Bx - C)$ ，其中相位移动的距离为 $\frac{C}{B}$ 。

2. **计算相位移动距离**：

- 对于 $y = 3 \cos(2x - \frac{\pi}{3})$ ， $B = 2$ ， $C = \frac{\pi}{3}$ 。
- 相位移动距离为 $\frac{C}{B} = \frac{\frac{\pi}{3}}{2} = \frac{\pi}{6}$ 。

3. **确定平移方向**：

- 如果函数形式为 $y = 3 \cos(2x - \frac{\pi}{3})$ ，即 C 为正，表示图象向右平移。
- 如果函数形式为 $y = 3 \cos(2x + \frac{\pi}{3})$ ，即 C 为负，表示图象向左平移。

4. **选择正确答案**：

- 根据题目描述，函数形式为 $y = 3 \cos(2x - \frac{\pi}{3})$ ，所以图象应向右平移 $\frac{\pi}{6}$ 个单位。


选项分析：

- **A**：右移 $\frac{\pi}{3}$ 个单位 —— 错误，距离计算不正确。
- **B**：左移 $\frac{\pi}{3}$ 个单位 —— 错误，方向和距离都不正确。
- **C**：右移 $\frac{\pi}{6}$ 个单位 —— 正确，方向和距离都正确。
- **D**：左移 $\frac{\pi}{6}$ 个单位 —— 错误，方向不正确。

最终答案：

C：右移 $\frac{\pi}{6}$ 个单位。

Image:



Question:

这题选啥?

Response:

这道选择题的题目是：

若 $(m = \sqrt[5]{(\pi - 3)^5})$, $(n = \sqrt[4]{(\pi - 4)^4})$, 则 $(m + n =)$ ()

选项：

- A. -7
- B. -1
- C. 1
- D. 7

解题步骤：

1. **计算 (m) 的值**：

- $(m = \sqrt[5]{(\pi - 3)^5})$
- 由于 $(\pi \approx 3.1416)$, 所以 $(\pi - 3 \approx 0.1416)$ (正数)。
- 奇数次根号的性质: $(\sqrt[5]{x^5} = x)$ (无论 (x) 是正数还是负数)。
- 因此, $(m = \pi - 3)$ 。

2. **计算 (n) 的值**：

- $(n = \sqrt[4]{(\pi - 4)^4})$
- $(\pi - 4 \approx -0.8584)$ (负数)。
- 偶数次根号的性质: $(\sqrt[4]{x^4} = |x|)$ (结果为非负数)。
- 因此, $(n = |\pi - 4| = 4 - \pi)$ 。

3. **计算 $(m + n)$ 的值**：

- $(m + n = (\pi - 3) + (4 - \pi))$
- 化简: $(\pi - 3 + 4 - \pi = 1)$ 。

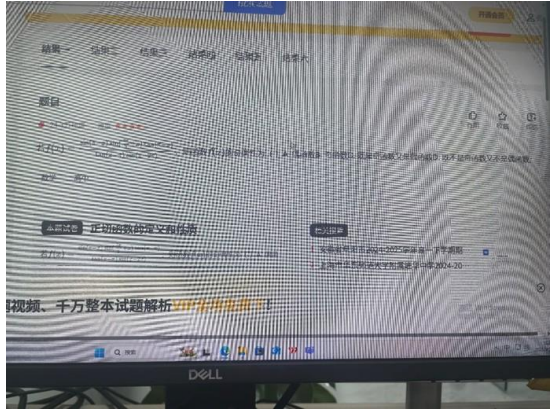
结论：

$(m + n = 1)$, 对应选项 **C**。

答案：

选 **C**。

Image:



Question:

帮我解下这道题

Response:

题目要求判断函数 $f(x) = \frac{\sin(\pi - x) \sin\left(\frac{3\pi}{2} + x\right) \tan(\pi - x)}{\tan(x - \pi) \sin(x - 2\pi)}$ 的奇偶性。以下是详细解答步骤：

步骤 1：简化函数表达式

利用三角函数的周期性、奇偶性和恒等变换性质，逐步简化分子和分母：

1. **分子部分**：

- $\sin(\pi - x) = \sin x$ (正弦函数的性质)
- $\sin\left(\frac{3\pi}{2} + x\right) = -\cos x$ (诱导公式)
- $\tan(\pi - x) = -\tan x$ (正切函数的性质)

分子简化为：
 $\sin x \cdot (-\cos x) \cdot (-\tan x) = \sin x \cos x \tan x$

2. **分母部分**：

- $\tan(x - \pi) = \tan x$ (正切函数的周期性)
- $\sin(x - 2\pi) = \sin x$ (正弦函数的周期性)

分母简化为：
 $\tan x \cdot \sin x$

3. **合并简化后的分子和分母**：

$$f(x) = \frac{\sin x \cos x \tan x}{\tan x \sin x} = \cos x$$

步骤 2：判断奇偶性

根据简化结果 $f(x) = \cos x$ ，分析其奇偶性：

- **偶函数定义**： $f(-x) = f(x)$
- **奇函数定义**： $f(-x) = -f(x)$

计算 $f(-x)$ ：
 $f(-x) = \cos(-x) = \cos x = f(x)$

因此， $f(x)$ 满足偶函数的定义。

步骤 3：验证选项

题目选项为：

- A. 偶函数
- B. 奇函数
- C. 既是奇函数又是偶函数
- D. 既不是奇函数又不是偶函数

正确答案是 **A**。

关键点总结

1. 通过三角恒等变换将复杂函数简化为基本形式 $\cos x$ 。
2. 偶函数的性质直接决定了答案。
3. 注意排除干扰项（如选项 C 要求 $f(x) \equiv 0$ ，但 $\cos x$ 不满足）。

最终答案： \boxed{A}