

3机12卡微调 Llama-2-7B 模型

a. 一共创建3个notebook

master

创建NoteBook ⓘ

* 名称

notebook-20240908-master

描述 ⓘ

请输入NoteBook描述

0/300

* 算法名称

LLM-xtuner ⓘ

* 算法版本

V1

* 镜像类型

预置镜像

* 镜像名称

aihpc2:v0.0.1

数据集类型

公共数据集

数据集名称

huggingface-cache

数据集版本

V1

* 资源池

second-train-pool

* 资源规格

24核-192G-4DCU(64GB)

自动停止

☐

[高级设置](#)

取消

确定

node0

创建NoteBook ⓘ

* 名称

notebook-20240908-node0

描述 ⓘ

请输入NoteBook描述

0/300

* 算法名称

LLM-xtuner ⓘ

* 算法版本

V1

* 镜像类型

预置镜像

* 镜像名称

aihpc2:v0.0.1

数据集类型

公共数据集

数据集名称

huggingface-cache

数据集版本

V1

* 资源池

second-train-pool

* 资源规格

24核-192G-4DCU(64GB)

自动停止

☐

[高级设置](#)

取消

确定

node1

创建NoteBook

×

* 名称

notebook-20240908-node1

描述

请输入NoteBook描述

0/300

* 算法名称

LLM-xtuner

ⓘ

* 算法版本

V1

▼

* 镜像类型

预置镜像

▼

* 镜像名称

aihpc2:v0.0.1

数据集类型

公共数据集

▼

数据集名称

huggingface-cache

数据集版本

V1

▼

* 资源池

second-train-pool

▼

* 资源规格

24核-192G-4DCU(64GB)

▼

自动停止

☐


高级设置

取消

确定

b. 打开三个notebook

名称	规格	描述	算法	镜像	状态	创建时间	操作
notebook-20240908-node1	24核-192G-4DCU(64GB显存)-1B8时/h		LLM-xtuner	aihpc2	运行中	2024-09-08 21:39:19	打开 停止 保存 详情
notebook-20240908-node0	24核-192G-4DCU(64GB显存)-1B8时/h		LLM-xtuner	aihpc2	运行中	2024-09-08 21:38:40	打开 停止 保存 详情
notebook-20240908-master	24核-192G-4DCU(64GB显存)-1B8时/h		LLM-xtuner	aihpc2	运行中	2024-09-08 21:37:55	打开 停止 保存 详情



c. 打开终端，查询master节点的ip

```

root@mee97dee8464411cb237a2355a493c60-task0-0:/# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1480
    inet 10.244.34.87 netmask 255.255.255.255 broadcast 10.244.34.87
    ether ba:4c:34:fd:51:18 txqueuelen 0 (Ethernet)
    RX packets 1655 bytes 397696 (397.6 KB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 1250 bytes 6284092 (6.2 MB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    loop txqueuelen 1000 (Local Loopback)
    RX packets 0 bytes 0 (0.0 B)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 0 bytes 0 (0.0 B)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

root@mee97dee8464411cb237a2355a493c60-task0-0:/# 

```

d. 在每个notebook的终端输入命令

a. 设置环境变量

```

# 设置环境变量
export NCCL_DEBUG=INFO
export NCCL_IB_DISABLE=0
export NCCL_IB_HCA=mlx5
export NCCL_SOCKET_IFNAME=eth0
export GLOO_SOCKET_IFNAME=eth0
export HF_HOME=/code/huggingface-cache/

```

b. 开始训练

```

# 开始训练
cd /code/
# 在 Notebook master 中执行
# 需要先在 Notebook0 中执行 ifconfig 查看 IP 地址, 比如这里得到
NPROC_PER_NODE=4 NNODES=3 PORT=12345 ADDR=10.244.34.87 \
NODE_RANK=0 xtuner \
train llama2_7b_chat_lora_lawyer_e3_copy.py \
--work-dir /userhome/xtuner-workdir \
--deepspeed deepspeed_zero3_offload
# 在 Notebook node0 中执行
NPROC_PER_NODE=4 NNODES=3 PORT=12345 ADDR=10.244.34.87 \
NODE_RANK=1 xtuner \
train llama2_7b_chat_lora_lawyer_e3_copy.py \
--work-dir /userhome/xtuner-workdir \
--deepspeed deepspeed_zero3_offload

```

```
# 在 Notebook node1 中执行
NPROC_PER_NODE=4 NNODES=3 PORT=12345 ADDR=10.244.34.87 \
NODE_RANK=2 xtuner \
train llama2_7b_chat_lora_lawyer_e3_copy.py \
--work-dir /userhome/xtuner-workdir \
--deepspeed deepspeed_zero3_offload
```

c. 在master节点查看运行情况

```
09/08 21:54:57 - mmengine - INFO - xtuner_dataset.timeout = 0:20:00
Generating train split: 81885 examples [00:00, 130138.72 examples/s]
Map (num_proc=32) - 100%
Filter (num_proc=32) - 100%
Map (num_proc=32) - 100%
Filter (num_proc=32) - 100%
Map (num_proc=32) - 100%
Filter (num_proc=32) - 100%
09/08 21:55:14 - mmengine - INFO - xtuner_dataset.timeout = 0:30:00
Generating train split: 92373 examples [00:00, 22733.29 examples/s]
Map (num_proc=32) - 100%
Filter (num_proc=32) - 100%
Map (num_proc=32) - 100%
Filter (num_proc=32) - 100%
Map (num_proc=32) - 100%
Filter (num_proc=32) - 100%
09/08 21:55:41 - mmengine - INFO - Dataset ConcatDataset has no meta info, dataset_meta in vizualizer will be None.
[2024-09-08 21:55:53.451] [INFO] [partition_parameters.py:345: _init_] Finished initializing model - num_params = 291, num_class = 8,748
Loading checkpoint shards: 100%
Loading checkpoint shards: 100%
Loading checkpoint shards: 100%
Loading checkpoint shards: 100%
09/08 21:56:47 - INFO [logging.py:96:log_dist] [Rank 0] DeepSpeed info: version=0.12.5, git=hash=7d4046, git-branch=HEAD
10908 21:56:48.337759 2620 ProcessGroupNCCL [Rank 0] ProcessGroupNCCL initialization options: NCCL_ASYNC_ERROR_HANDLING: 1, NCCL_DEBUG: 0, NCCL_ENABLE_TIMING: 0, NCCL_BLOCKING_WAIT: 0, TIMEOUT(ms): 1800000, USE_HIGH_PRIORITY_STREAM: 0, TORC
6_DISABLE_DEBUG: 0, NCCL_DEBUG: INFO, ID=41981958
[2024-09-08 21:56:48.358] [INFO] [logging.py:96:log_dist] [Rank 0] DeepSpeed Flops Profiler Enabled: False

steps_per_print : 1.000000e+12
}
09/08 21:56:52 - mmengine - INFO - Num train samples 143958
09/08 21:56:52 - mmengine - INFO - train example:
09/08 21:56:52 - mmengine - INFO - <> [INST] <<SYS>>
You are a helpful, respectful and honest assistant. Always answer as helpfully as possible, while being safe. Your answers should not include any harmful, unethical, racist, sexist, toxic, dangerous, or illegal content. Please ensure that your responses are socially unbiased and positive in nature.
你是一名专业的中国律师，请根据用户的问题给出准确、有理有据的回复。
<</SYS>>
[/INST] [INST] 你好，我昨晚刚租了房子，签了房屋租赁合同，付一押一，合同签的是今年9.9 [/INST] 年9月到明年8月底，租期一年。</s>
09/08 21:56:52 - mmengine - INFO - before_train in EvaluateChatHook.
[2024-09-08 21:56:58.861] [WARNING] [parameter_offload.py:35: _split_to_tensors_only] A module has unknown inputs or outputs type ((class 'transformers.cache_utils.DynamicCache')) and the tensors embedded in it cannot be detected. The FaSO-3 hooks designed to trigger before or after backward pass of the module relies on knowing the input and output tensors and therefore may not get triggered properly.
[]

10908 22:04:52.379523 2620 ProcessGroupNCCL [Rank 0] NCCL_DEBUG: INFO
09/08 22:04:52 - mmengine - INFO - Iter(train) [ 1/35991] lr: 2.0000e-09 eta: 8 days, 13:59:51 time: 20.6055 data_time: 0.0235 memory: 1518 loss: 1.0194
09/08 22:05:04 - mmengine - INFO - Iter(train) [ 2/35991] lr: 1.8753e-07 eta: 6 days, 21:12:03 time: 16.5252 data_time: 0.0285 memory: 3600 loss: 1.1006
09/08 22:05:23 - mmengine - INFO - Iter(train) [ 3/35991] lr: 3.7305e-07 eta: 7 days, 2:56:48 time: 17.1004 data_time: 0.0272 memory: 3547 loss: 1.1680
09/08 22:05:36 - mmengine - INFO - Iter(train) [ 4/35991] lr: 5.5853e-07 eta: 6 days, 17:00:17 time: 16.1063 data_time: 0.0257 memory: 3568 loss: 1.2011
09/08 22:05:50 - mmengine - INFO - Iter(train) [ 5/35991] lr: 7.4411e-07 eta: 6 days, 13:55:44 time: 15.7990 data_time: 0.0253 memory: 3560 loss: 1.2695
09/08 22:06:05 - mmengine - INFO - Iter(train) [ 6/35991] lr: 9.2963e-07 eta: 6 days, 12:23:33 time: 15.6458 data_time: 0.0253 memory: 3645 loss: 1.2814
09/08 22:06:18 - mmengine - INFO - Iter(train) [ 7/35991] lr: 1.1152e-06 eta: 6 days, 8:41:49 time: 15.2765 data_time: 0.0252 memory: 3607 loss: 1.2344
[]
```