



# ??LLM?? ????????

??

1. [ ]
2. [ ]
3. [ ]
  1. [ ]
  2. [ ]

??

1. [ ]
  1. [ ] , [ ] , [ ] , [ ]
  2. [ ] , [ ] , [ ] , [ ] , [ ]
  3. [ ]
  4. [ ] , [ ] , [ ] , [ ] , [ ] , [ ]
2. [ ]
  1. [ ]
  2. [ ] AI [ ]
3. [ ]
  1. [ ] AI
4. [ ] 1 [ ] token [ ]
  1. [ ] Yes [ ] No [ ] 996563 [ ] 365336
  2. [ ] Yes [ ] No [ ] 9.11 > 9.9
5. [ ]
  1. Anthropic [ ]  
[https://github.com/anthropics/original\\_performance\\_takehome/blob/main/problem.py](https://github.com/anthropics/original_performance_takehome/blob/main/problem.py)
  2. [ ] [https://github.com/deepreinforce-ai/IterX-tutorials/tree/main/anthropic\\_take\\_home](https://github.com/deepreinforce-ai/IterX-tutorials/tree/main/anthropic_take_home) [ ] Anthropic [ ]
6. [ ]
  1. [ ]
  2. [ ]

# ???????

```

Mops , K
stride pooling
16000
1D
```

```

conv_dim: Tuple[int, ...] = (8, 16, 32, 64, 128, 256, 16) # Progressive increase to 512
conv_kernel: Tuple[int, ...] = (10, 3, 3, 3, 3, 2, 2) # First layer larger kernel
conv_stride: Tuple[int, ...] = (5, 2, 2, 2, 2, 2, 2) # First layer larger stride
conv_bias: Tuple[bool, ...] = (False, False, False, False, False, False, False)
conv_padding: Tuple[int, ...] = (0, 0, 0, 0, 0, 0, 0)
# Disable depthwise for stability
conv_use_depthwise: Tuple[bool, ...] = (False, False, False, False, False, False, False)
# Reduce pooling - only 2 layers
conv_use_pool: Tuple[bool, ...] = (False, False, False, False, False, False, False)
# Disable residual for now (can cause gradient issues)
conv_use_residual: Tuple[bool, ...] = (False, False, False, False, False, False, False)
conv_pool_type: Tuple[str, ...] = ("avg", "avg", "avg", "avg", "avg")
conv_pool_kernel: Tuple[int, ...] = (2, 2, 2, 2, 2)
conv_pool_stride: Tuple[int, ...] = (2, 2, 2, 2, 2)
conv_pool_padding: Tuple[int, ...] = (0, 0, 0, 0, 0)
```

# ???????

1. 20260228 106.45K 136.01Mops
2. Hunyuan 20260228 2440 K 3722 MOPs
3. DeepseekV3.2 98,448 ~1.59 MOPs
4. **GPT-5.2 106,448 16,324,528MACs 32.649 Mops**
5. **GLM-5 16.33 Mops 106.45 K**
6. Qwen3.5-Plus 139.2 K (139,216 ) 19.7 M (19,712,000 )
7. kimi k2.5 106.688 K (107K) 10.99 Mops (11 Mops)
8. Claude-Haiku-4.5 ~106.4K ~620.5 MOPs 32
9. **Claude-Sonnet-4.5 106.4K 16.33 MOPs**
10. **Claude-Opus-4.5 106.45 K 16.33 MOPs**

# ??

1. **Gemini-3-Flash 106.45 K 32.66 Mops**
2. DeepseekV3.2 ~105.4 K ~58.1 Mops

3. 20260228 106.45 K 16.33 Mops

4. Minimax-M2.5 106448 104K 16328752 16.33M

5. GLM5 106.45 K (106,448 ) 16.33 M

??????????

# Test1

git : `http://deve.work:33333/colin/mde.git`

commit : `c61827061d6a34cbd2ca2fae62b59fcb780f192e`

prompt : `vscode markdown`

fix commit : `3e64dfab0643fecc583bb1e77467373d7e43854e`

1. 1

1. Lute (3.8MB)

1. `<script>` `onload` `id="vditorLuteScript"` `Vditor` `addScript()` `lute` `id` `script` `3.8MB`

2. `window.__luteReady` `promise` `lute` `_boot()` `lute` `id` `Vditor`

2. Icons XHR (43KB)

1. `<script defer>` `ant.js` `id="vditorIconScript"` `Vditor` `addScriptSync()` `id` `XHR` `---`

2. `defer` `script` `id="vditorIconScript"` `addScriptSync`

3. Loading 1200ms

1. `Vditor ready` 1200ms `loading`

2. `ready`

1. 1

1. Opus4.7 500K

# Test2

git : `http://deve.work:33333/colin/mde.git`

commit : `5b9134576b7a11b52dc1c7962d2827c487c698bc`

prompt : `inline code` `inling code` `inling code`



drawImage GPU canvas lane lane  
BGRA  
lane

GPU GPU  
v7.21

canvas drawImage lane ImageData putImageData canvas CPU  
lane  
6 × 5 0 JS

GPU canvas canvas

stage



```
_____ |
| indexer | v7 pass _____
```

```
| .py
```

```
|
```

```
_____ |
```

```
_____
```

```
_____ |
```

```
_____
```

```
- _____ 350ms _____4 _____ HEAD
```

```
_____
```

```
- src/utils/insight/ _____
```

```
JSON[1.png]2.png[time*.png]insight_sentinel_*.json_____
```

```
commit: 20a88b48fde3fdf1bcc09659f65df6da5deb3de3
```