Binary AI

二值 二进制 Binary 神经网络算法
先量化 再训练
语言模型究竟能记住多少内容？

?? ??? Binary ??????

??

1. FPGA[illegible] LUT([illegible])[illegible] LUT[illegible]
    1. [illegible]
    2. [illegible]
2. FPGA[illegible] AI[illegible]
    1. [illegible] PPA
    2. [illegible] FPGA[illegible]
    3. [illegible] PPA[illegible]
3. [illegible] FPGA[illegible] LUT[illegible] AI[illegible]
    1. [illegible]
4. [illegible]
    1. [illegible]
    2. [illegible] LUT[illegible]
5. binary LUT [illegible] bit[illegible]
    1. [illegible] 8bit[illegible] 256bit[illegible]
    2. [illegible]
    3. [illegible] CIM[illegible]
6. [illegible] Bitnet[illegible]
    1. [illegible] FP8[illegible] int4[illegible]
    2. [illegible]

??

1. LUT6
    1. 6[illegible] 1[illegible] LUT[illegible] 64[illegible] 64bit
    2. xc7k480t[illegible] 74659 slices[illegible] slice[illegible] 4 LUT[illegible] 8[illegible] flip-flop
    3. [illegible] **lut**[illegible] 64bit*74659*4 = 2389088 byte[illegible] 74659*4 / 6 * 2 * 0.5G [illegible] GFlops[illegible] = 50 TFlops
    4. [illegible] LUT[illegible] LUT[illegible]
2. [illegible]
    1. [illegible]
    2. 4[illegible] 6LUT[illegible] 8bit lut[illegible] 8[illegible] LUT8 [illegible] 8bit[illegible] 8[illegible] 8[illegible]
    3. [illegible]
3. [illegible] LUT[illegible]
    1. [illegible] LUT[illegible]
        1. “[illegible]+weight”[illegible] LUT

        2. □□□□□□□□□□ LUT□□□ LUT
           □□□
           LUT□□□□□□□□□□□□□□□□□ LUT□□□□
            1. □□□□□□□□□□□□ GPU□□□□□□□□
            2. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
            3. □□□□□□□□ AI□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
               FP8□□□□□□□□
    2. □□□ □□□□ □□□□□□□□□□□□□
4. □□□□□□ LUT□□□□□□□□□ f(x)=y□□□□□□□□□□□ LUT□□
    1. ElementWise□□ CNN□ GEMM□□□□ weight□□□□ LUT□□
        1. □□□□□□□□□□□□□□ LUT□□
    2. Reduction□ normal□ softmax□□□□□□ Reduce□□
    3. □□ GEMM□□□□□□□□□□□□□□□□□□□□□□□□□
5. □□ LUT
    1. LUT□□□□□□□□□□□□□□□□□□□□□□□ FPGA□□□□□□□□□□□□

?????

????

1. □□□□□□□□ "□□ ”□□□□□□□□□
2. □□□□□□□□□ softmax□ normal□□□□□□□□□□ GEMM□ CNN□□□□□□□□□□□
3. □□□□□□□□□□□□□□□□□□□□□□□□□□□ LUT□□□□□□□□□□□□□□
4. □□
   □

??????

1. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□ LUT□□□□□□□□□□□□□□
    1. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
    2. □□□□□□□□ ”□□□□ “
       □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
    3. □□□
       □□□□□□□□□□□□□□□□□□□□□□□□
2. □□□□□□□□
    1. □□□□□□□□□□□□□□□□□□□□□
    2. □□□□□□□□□□□□□□□
        1. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□
    3. □□□□□□□□□□□□□□□□□□
        1. □□□ bit□□□□□□□□ XOR□□□□□□□□□□□□□□ AND□□
    4. □□ LUT
       □□□
3. □□□
    1. □□□□□□□□□□□□□□□□□□□ linear
        1. □□□□□□□□□ =□□□□□□□□□□□ =□□□□

    2. “ +weight” LUT
  2. =>

AI???????

1.
2.
   1.
3.
4.
5.
6.
7. shift scale

LLM??

1. ” “ bit
   1.
2.
   1.
   2. ” “
   3.
3. 8bit
   1. bit
   2. 8bit 256 weight
   3. weight bit
      1.
   4.
      1. LUT
      2. LUT bit LUT
      3.
4.
   1. /
   2. “ ” “ ”
   3.
      1.
      2.
      3.
5.
   1. x -> x
      1. attention qkv bit
         1. QK qk input
   2. LUT reduce
6. **LUT**
   1.

??

1. [illegible] LLM/CNN
    1. [illegible]
    2. [illegible] PPA
    3. [illegible] FPGA[illegible]
2. [illegible]

???

1. [illegible]
2. AI[illegible]
3. [illegible]
4. [illegible] DNA[illegible]
5. [illegible] [illegible]
6. [illegible] LUT[illegible]
    1. [illegible]
    2. [illegible]
    3. [illegible] scale[illegible]

??????

1. BitNet b1.58 https://mp.weixin.qq.com/s/G9ZbMnBVbeH1m45HY2JIKA
2. [illegible]

??

1. [illegible]
2. [illegible]

??

1. https://arxiv.org/html/2502.19008v1
2. [illegible] Binary Neural Networks[illegible] https://zhuanlan.zhihu.com/p/117285043
3. https://github.com/ryuz/BinaryBrain
4. [illegible] https://www.slideshare.net/kentotajiri/ss-77136469
5. BinaryConnect: Training Deep Neural Networks with binary weights during propagations https://arxiv.org/pdf/1511.00363.pdf
6. Binarized Neural Networks https://arxiv.org/abs/1602.02505
7. Binarized Neural Networks: Training Deep Neural Networks with Weights and Activations Constrained to +1 or -1 https://arxiv.org/abs/1602.02830
8. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks https://arxiv.org/abs/1603.05279

9. Xilinx UltraScale Architecture Configurable Logic Block User Guide https://japan.xilinx.com/support/documentation/user_guides/ug574-ultrascale-clb.pdf
10. 1-bit AI Infra: Part 1.1, Fast and Lossless BitNet b1.58 Inference on CPUs https://arxiv.org/abs/2410.16144v2
11. Bitnet.cpp: Efficient Edge Inference for Ternary LLMshttps://arxiv.org/abs/2502.11880v1
12. Continual Quantization-Aware Pre-Training: When to transition from 16-bit to 1.58-bit pre-training for BitNet language models?https://arxiv.org/abs/2502.11895v1
13. (NEW!) BitNet v2: Native 4-bit Activations with Hadamard Transformation for 1-bit LLMs https://arxiv.org/abs/2504.18415
14. (NEW!) BitVLA: 1-bit Vision-Language-Action Models for Robotics Manipulation https://arxiv.org/abs/2506.07530

??? ???

?????LLM?Binary

1. Lut□□□□□□□□□□□□□□□□
    1. □□□□　　touch□□□□□□□□□□□□
    2. □□□□□□□□□□□□□　　Lut□□□

?????Bianry??

1. □□□□□
    1. □□□□□□□□□□

??????????????

???????How much do language models memorize?

□□□□□□□

- □□ □ How much do language models memorize?□□□□□□□□□□□□□□□□
- □□ □ John X. Morris □□□ Meta FAIR, Google DeepMind, NVIDIA □□□□□□□□□□□□□
- □□□□ □□□□□□□ LLM□□□□□□□□□□□□□□□ “□□□ Grokking□ ”□□□

1. ??????? (Motivation)

- □□□□ □□□□□□□□□□□□□□□□□□□□ “□□□□□ Memorization□ ”□ “□□□□□□ Generalization□ ”□□□□□
- □□□□ □□□□□□□□□□□□□□□□□□□□□□□□ “□□ ” □□□□□□□□□□□□□□□□□□□□□□□□□□□□

2. ??????? (Methodology)

□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□

1. □□□□□ **(Unintended memorization)** □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ “□□□□ Lookup Table □ ”□□□□□□□□□□□□□□□
2. □□ **(Generalization)**□□□□□□□□□□□□□□□□□□□□□□□□□□□□

□□□□ □

□□□□□□□□□□□□□□ □□□□□ “□□ ”□□□ □□□□□□□□□□□□□□□□□□□□ “□□□□ ”□□□□□□□□□□□□□□□□□□□□□□□□□□□□

3. ???? (Experiments)

- [illegible] [illegible] Transformer [illegible] **500K[illegible] 50[illegible] 1.5B[illegible] 15[illegible]** [illegible]
- [illegible] [illegible] “[illegible]”[illegible]“[illegible]” [illegible]

4. ??????? (Key Findings)

- [illegible] [illegible] GPT [illegible] Transformer [illegible] **3.6 bits / [illegible] 3.6 bits per parameter[illegible]** [illegible]
- [illegible] **“[illegible] (Grokking)”[illegible]** [illegible]
  - [illegible] “[illegible]”[illegible]
  - [illegible] [illegible] **3.6 bits/[illegible]** [illegible] [illegible]
  - [illegible] [illegible] **“[illegible] Grokking[illegible]”** [illegible] [illegible]
- [illegible] **(Scaling Laws)**[illegible] Scaling Laws[illegible] “[illegible]” [illegible]“[illegible]”[illegible] “[illegible] Membership Inference [illegible]”[illegible]

5. ????? (Takeaways)

[illegible] “[illegible]”[illegible] ——“[illegible] 3.6 bits[illegible]”[illegible] [illegible] [illegible] “[illegible]”[illegible] “[illegible]” [illegible] [illegible]