Algorithm

- OP frequency
  - ONNX OP statistics
- Model Compression
  - Model Quantization
  - Neural Arch Search
  - Knowledge Distillation
  - Pruning
  - reference
- Data type
  - data type

OP frequency

ONNX OP statistics

alphabet order

```
   3 Abs
4325 Add
  20 And
   2 ArgMax
 101 AveragePool
3467 BatchNormalization
3152 Cast
   8 CategoryMapper
  27 Ceil
 407 Clip
   2 Compress
2706 Concat
7714 Constant
 788 ConstantOfShape
6448 Conv
   2 ConvInteger
   4 ConvTranspose
   1 CumSum
3227 DequantizeLinear
 503 Div
 100 Dropout
  82 DynamicQuantizeLinear
   5 DynamicQuantizeLSTM
  97 Equal
  24 Erf
 194 Exp
 132 Expand
  40 Flatten
  52 Floor
```

12 FusedMatMul
7681 Gather
251 Gemm
29 GlobalAveragePool
8 Greater
2 Hardmax
91 Identity
150 InstanceNormalization
228 LeakyRelu
52 Less
1 LessOrEqual
91 Log
30 Loop
90 LRN
5 LSTM
861 MatMul
96 MatMulInteger
2 Max
522 MaxPool
4 Min
3536 Mul
3 Neg
688 NonMaxSuppression
843 NonZero
41 Not
4 OneHot
160 Pad
234 Pow
100 PRelu
1 preprocess
276 QLinearAdd
5 QLinearAveragePool
67 QLinearConcat
1191 QLinearConv
4 QLinearGlobalAveragePool
72 QLinearLeakyRelu
30 QLinearMatMul
62 QLinearMul
10 QLinearSigmoid

1577 QuantizeLinear
7 Range
100 Reciprocal
6 ReduceMax
466 ReduceMean
65 ReduceMin
8 ReduceSum
4794 Relu
2397 Reshape
43 Resize
48 RoiAlign
2 Round
2 Scan
12 Scatter
36 ScatterElements
2229 Shape
90 Sigmoid
2276 Slice
230 Softmax
135 Split
274 Sqrt
1700 Squeeze
879 Sub
249 Sum
146 Tanh
47 Tile
60 TopK
1469 Transpose
7973 Unsqueeze
22 Upsample
15 Where

Frequency order

1 CumSum
1 LessOrEqual
1 preprocess
2 ArgMax
2 Compress

2 ConvInteger
2 Hardmax
2 Max
2 Round
2 Scan
3 Abs
3 Neg
4 ConvTranspose
4 Min
4 OneHot
4 QLinearGlobalAveragePool
5 DynamicQuantizeLSTM
5 LSTM
5 QLinearAveragePool
6 ReduceMax
7 Range
8 CategoryMapper
8 Greater
8 ReduceSum
10 QLinearSigmoid
12 FusedMatMul
12 Scatter
15 Where
20 And
22 Upsample
24 Erf
27 Ceil
29 GlobalAveragePool
30 Loop
30 QLinearMatMul
36 ScatterElements
40 Flatten
41 Not
43 Resize
47 Tile
48 RoiAlign
52 Floor
52 Less
60 TopK
62 QLinearMul
65 ReduceMin
67 QLinearConcat
72 QLinearLeakyRelu
82 DynamicQuantizeLinear
90 LRN
90 Sigmoid
91 Identity
91 Log
96 MatMulInteger
97 Equal
100 Dropout
100 PRelu
100 Reciprocal

101 AveragePool
132 Expand
135 Split
146 Tanh
150 InstanceNormalization
160 Pad
194 Exp
228 LeakyRelu
230 Softmax
234 Pow
249 Sum
251 Gemm
274 Sqrt
276 QLinearAdd
407 Clip
466 ReduceMean
503 Div
522 MaxPool
688 NonMaxSuppression
788 ConstantOfShape
843 NonZero
861 MatMul
879 Sub
1191 QLinearConv
1469 Transpose
1577 QuantizeLinear
1700 Squeeze
2229 Shape
2276 Slice
2397 Reshape
2706 Concat
3152 Cast
3227 DequantizeLinear
3467 BatchNormalization
3536 Mul
4325 Add
4794 Relu
6448 Conv
7681 Gather
7714 Constant
7973 Unsqueeze

Model Compression

Model compression techniques are mainly classified into quantization (Zafrir et al.,2019), pruning (Hoefler et al., 2021; Vadera and Ameen, 2022), Neural Architecture Search(NAS) (Sun et al., 2019), and knowledge distillation

Model Quantization

Quantization Granularity

Quantization is a magic spell to reduce the memory footprint of a model. But often quantization leads to a drop in the accuracy of the model. This is where the Granularity of quantization comes into the picture. Selecting the right granularity helps in maximizing the quantization without much drop in the accuracy performance

Per-Tensor Quantization

In per-tensor quantization, the same quantization parameters are applied to all elements within a tensor. Applying the same parameters across tensors will cause a drop in accuracy.

Per-Channel Quantization

In per-channel quantization, different quantization parameters are applied to each channel of a tensor independently. This often leads to a lower error while quantizing compared to per tensor quantization.

Per-channel quantization captures variations in different channels more accurately. This usually helps in CNN models where the range of weights varies over different channels.

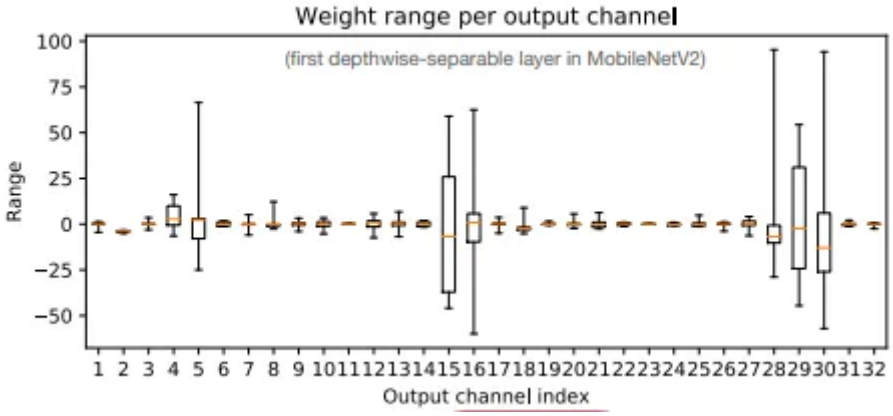

Mix-precision Quantization

different layer using different precision.

Quantization Method

Quantization Method	Accuracy	Model Size Reduction	Inference Speed	Ease of Implementation
GGUF	High (adaptive quantization of parameter groups)	Significant (grouped quantization)	High (optimized through grouped quantization)	Moderate (requires parameter grouping)
AWQ	High (asymmetric quantization for accuracy)	Moderate	Moderate to High (asymmetric approach)	Moderate (asymmetric scaling)
GPTQ	Moderate (post-training fine-tuning)	Moderate	Moderate (benefits from post-quantization)	Easy (applied post-training)
GGML	High (mixed-precision within groups)	Significant (effective group-wise precision)	High (mixed-precision within groups)	Complex (group-wise mixed-precision)
PTQ	Moderate (quick and easy implementation)	High (straightforward implementation)	High (quick deployment)	Easy (quick application)
QAT	Very High (trained with quantization in mind)	Moderate to High (requires complex training)	High (optimized for quantization during training)	Complex (requires training with quantization)
Dynamic Quantization	Moderate (applied during inference)	High (applied at inference time)	Moderate to High (quick deployment)	Easy (inference-based)
Mixed-Precision	High (selectively applied different precisions)	Moderate to High (selectively applied)	High (selective precision boosts performance)	Complex (selective precision assignment)

Neural Arch Search

Knowledge Distillation

Pruning

reference

[illegible]

[illegible] CNN[illegible]

[illegible] CNN[illegible] INT8[illegible]

[illegible] (feature map) [illegible] 8-bit[illegible] 32-bit [illegible]

[illegible] (GEMM

)
[illegible] INT8[illegible]

CNN[illegible] CV [illegible]

[illegible] (LLM)[illegible] transformer

[illegible] (GPT)[illegible] GEMM[illegible] LLM[illegible]
CNN[illegible] (CNN[illegible] [illegible] white paper on

Quantization[illegible]

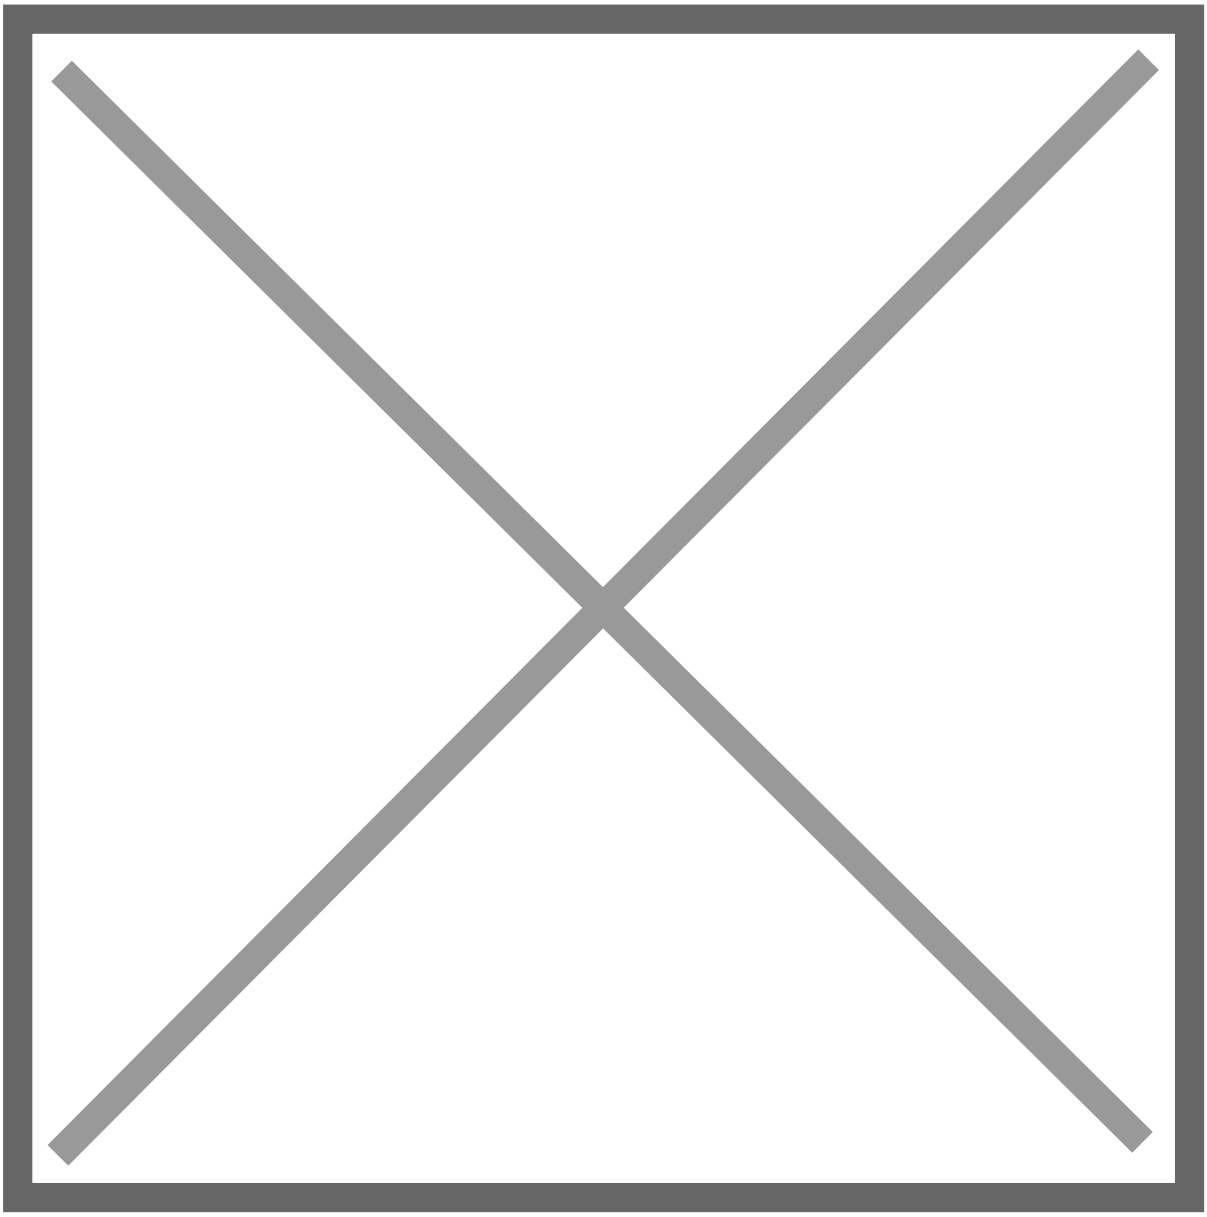

量化公式

这里主要关注最常见的对称量化，简单来说，INT8 量化即将浮点数通过缩放因子映射到范围在 [-128, 127] 内的 8bit 表示，量化公式如下:

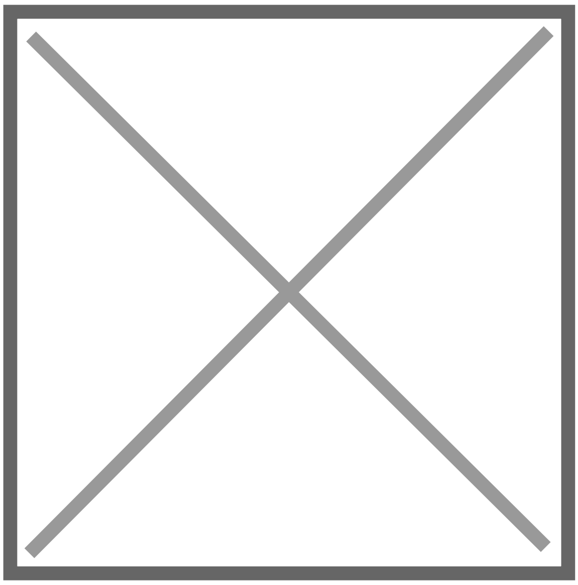

Round [illegible] Clip [illegible] (Outlier) [illegible] [-128, 127] [illegible] scale [illegible]

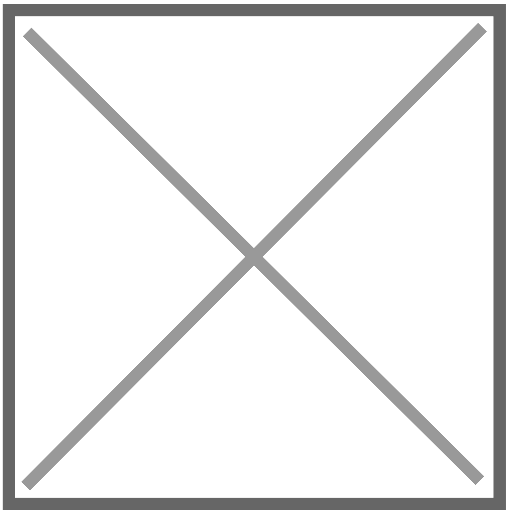

[illegible] Round[illegible] Clip [illegible] LLM[illegible] CNN[illegible] ([illegible]) [illegible] trade-off[illegible]

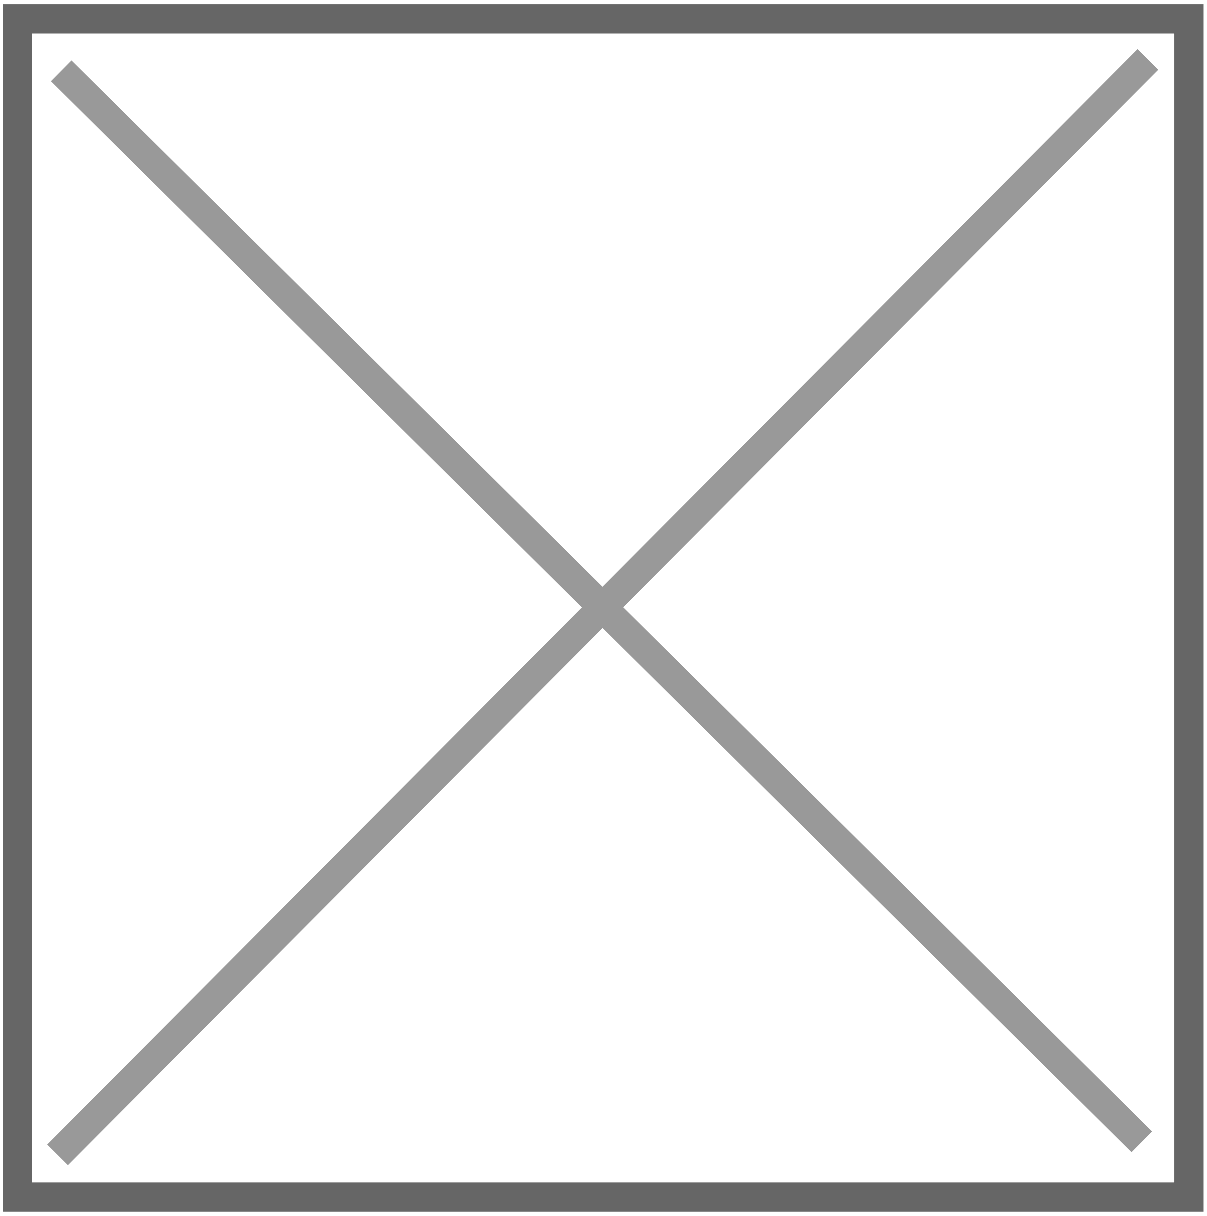

[illegible]

[illegible] tensor [illegible] scale [illegible] layer [illegible] tensor [illegible] scale [illegible] tensor [illegible] tensor [illegible] scale [illegible]

[illegible] CNN [illegible] tensor [illegible] per-tensor [illegible] tensor [illegible] ([illegible] tensor) [illegible] scale [illegible] per-channel [illegible]

[illegible]

[illegible] LLM [illegible] LLM [illegible] per-vector [illegible] per-group [illegible] vector [illegible]

[illegible]

[illegible] “[illegible]” [illegible] float32[illegible] float16 [illegible] bfloat16 [illegible] INT8[illegible] INT4[illegible]

[illegible] LLM[illegible]

- [illegible] `sign` [illegible] s[illegible] 1 bit[illegible]
- [illegible] `exponent` [illegible] k[illegible] 8 bits[illegible] 2 [illegible]
- [illegible] `fraction` [illegible] M[illegible] 23 bits[illegible]

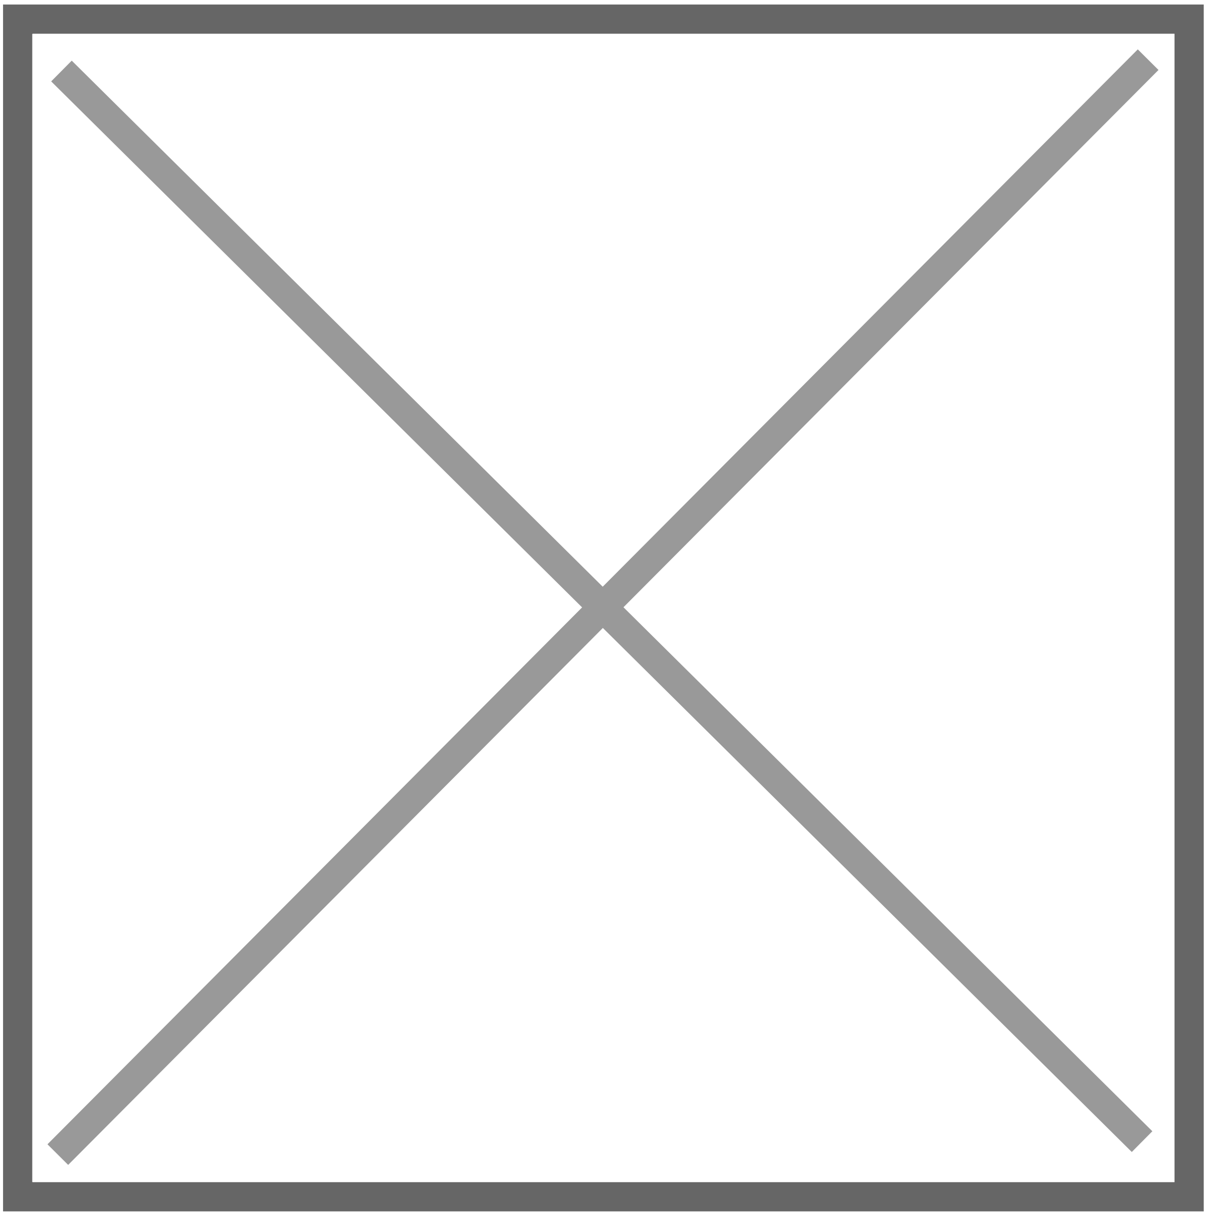

●**Float32 (FP32)**[illegible] IEEE 32 [illegible] 8 [illegible] 23 [illegible] 1 [illegible] FP32 [illegible]

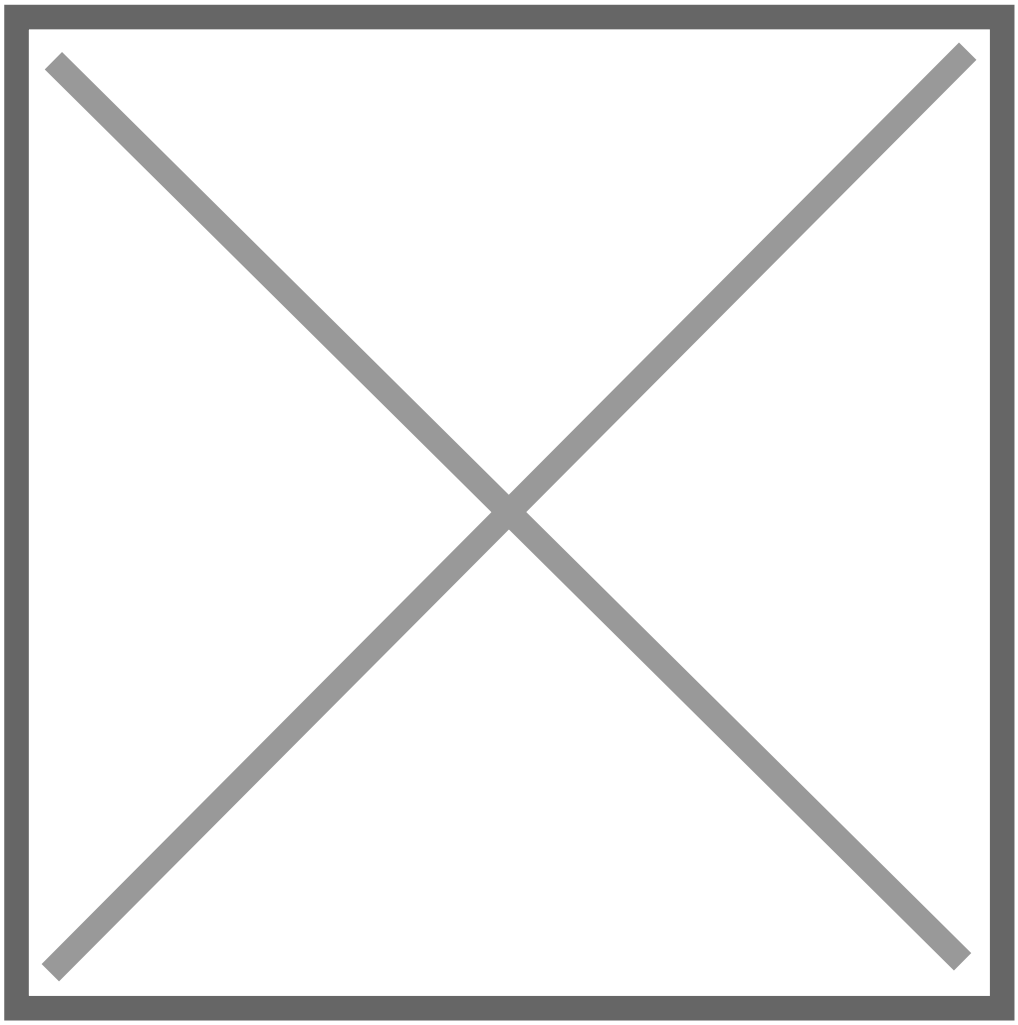

●**TensorFloat-32(TF32)** 。使用 19 位表示，结合了 BF16 的范围和 FP16 的精度，是计算数据类型而不是存储数据类型。目前使用范围较小。

●**Float16 (FP16)** 。指数 5 位，尾数 10 位，符号 1 位。 FP16 数字的数值范围远低于 FP32，存在上溢 (当用于表示非常大的数时) 和下溢 (当用于表示非常小的数时) 的风险，fp16混合精度训练中通过缩放损失 (loss scaling) 来缓解这个问题。

●**Bfloat16 (BF16)**。指数 8 位 (与 FP32 相同)，尾数 7 位，符号 1 位。这意味着 BF16 可以保留与 FP32 相同的动态范围。但是相对于 FP16，损失了 3 位精度。因此，在使用 BF16 精度时，大数值绝对没有问题，但是精度会比 FP16 差。

●**(Un)signed Char (INT8)**。定点数 7位，符号 1位。用来表示 [-128, 127] 之间的整数。根据量化方式不同，有时也会转而用 UINT8来表示 [0, 255]之间的整数。

●**Unsigned Char (INT4 * 2)**。一个字节的 Unsigned Char可以表示（装入）两个 INT4的值。定点数数 3 位，符号 1位。用来表示 [-8, 7]之间的整数。

CNN量化 vs. LLM量化

LLM[illegible] CNN[illegible]

[illegible] [illegible] CNN[illegible] per-channel[illegible] per-layer[illegible] kernel[illegible] LLM [illegible] vector[illegible] (tensor[illegible]) [illegible] A*B=C[illegible] per-tensor[illegible] A [illegible] B [illegible] INT [illegible] per-vector[illegible] LLM[illegible] per-vector[illegible] LLM [illegible] X[illegible] W [illegible] FP16[illegible] k[illegible] k [illegible] 2[illegible] k[illegible] 256[illegible] 128[illegible] per-group[illegible]

[illegible]

[illegible] ***INT4***[illegible] CNN[illegible] INT8[illegible] LLM [illegible] CNN[illegible] INT4 [illegible] 13B[illegible] scale [illegible] scale[illegible] fp16[illegible] fp32[illegible] scale[illegible] INT4 [illegible]

[illegible] [illegible] CNN[illegible] loss [illegible] per-channel [illegible] scale [illegible] LLM[illegible] head [illegible] tensor[illegible] hidden dim [illegible] [50, 100] [illegible] [0.5, 1.2], [illegible] LLM[illegible] attention[illegible] softmax[illegible] Attention Is Off By One[illegible] tensor[illegible] per-tensor[illegible] LLM [illegible]

[illegible] [illegible] CNN[illegible] LLM[illegible] cuda[illegible] INT [illegible] Core [illegible]
[illegible]
[illegible] [illegible]

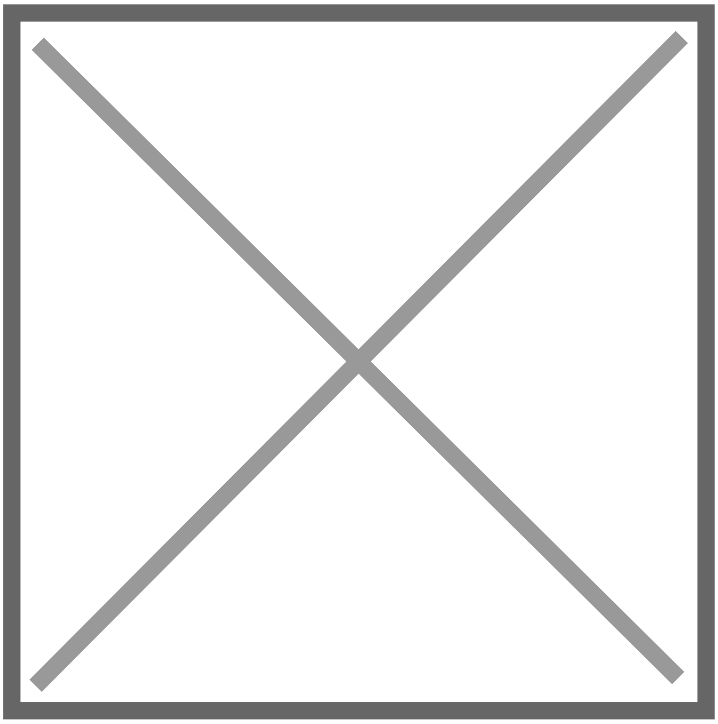

LLM

LLM

LLM.int8():

bit

CNN PTQ (KL)

CNN

LLM tensor

6B layer

sequence length

quantization

hidden size

sequence token

paper

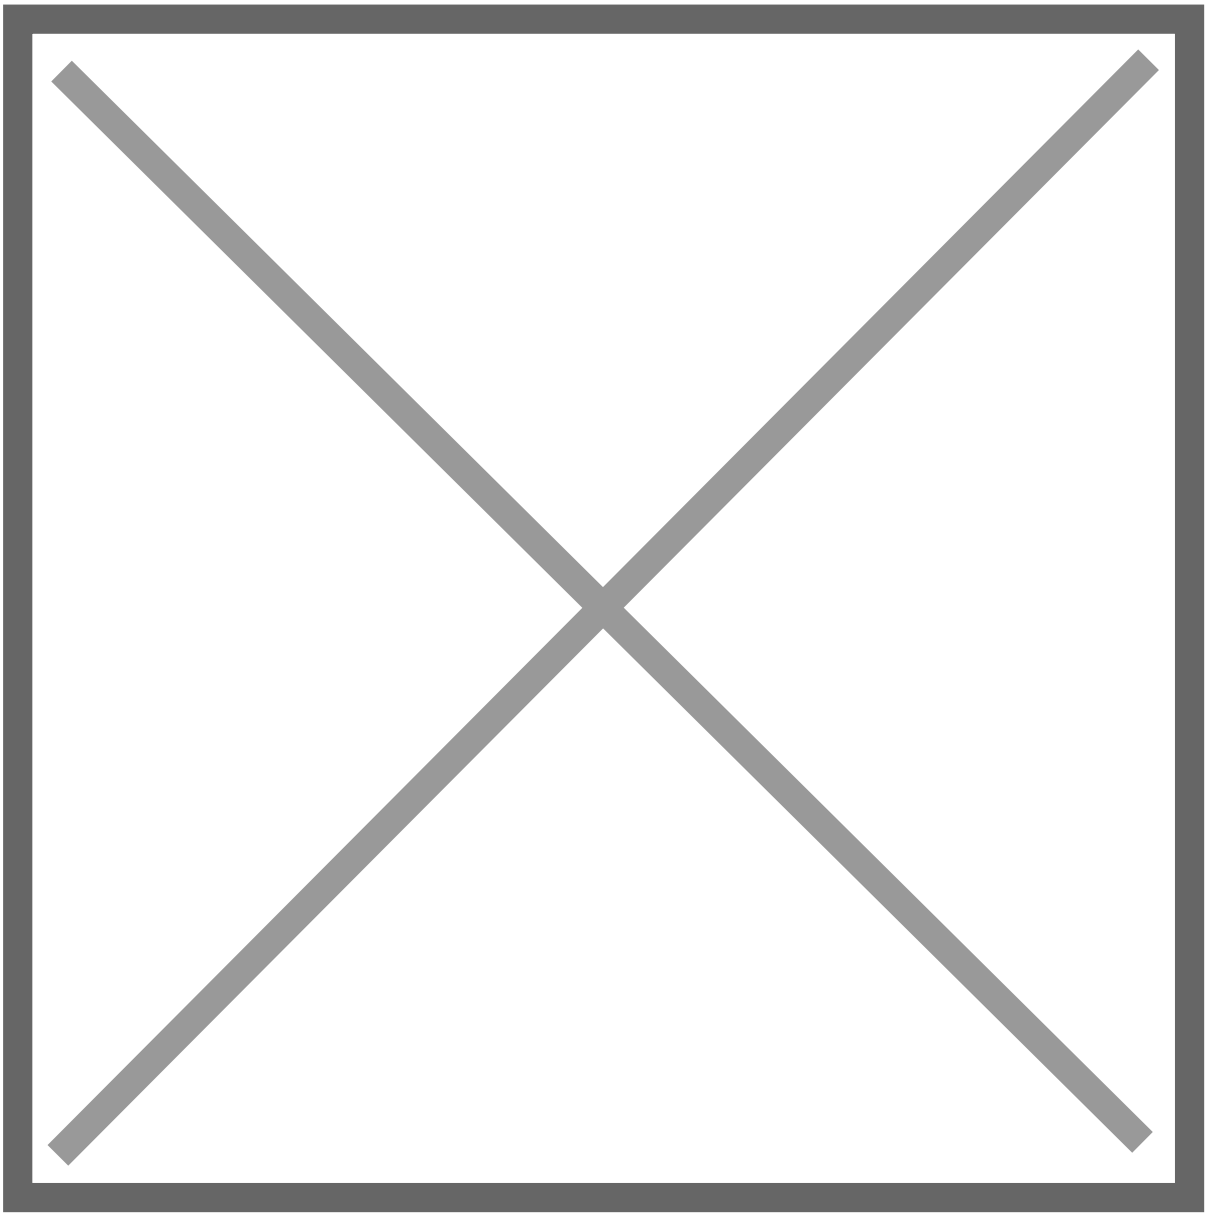

X activations sequence length hidden-size

X W int8 fp16 **(W token) element-wise concat**

LLM.INT8() gemm X W threshold 6 6

LLM.int8() LLM.int8() BLOOM-176B FP16 15% 23% (T5-3B T5-11B)

[illegible] LLM [illegible] block [illegible] block [illegible] fp16 [illegible] fp16 [illegible]

GPTQ: [illegible]

GPTQ [illegible] OBQ [illegible] GPTQ [illegible] OBQ [illegible] OBS [illegible] ([illegible] OBS [illegible] W_q , [illegible] δ_q [illegible]

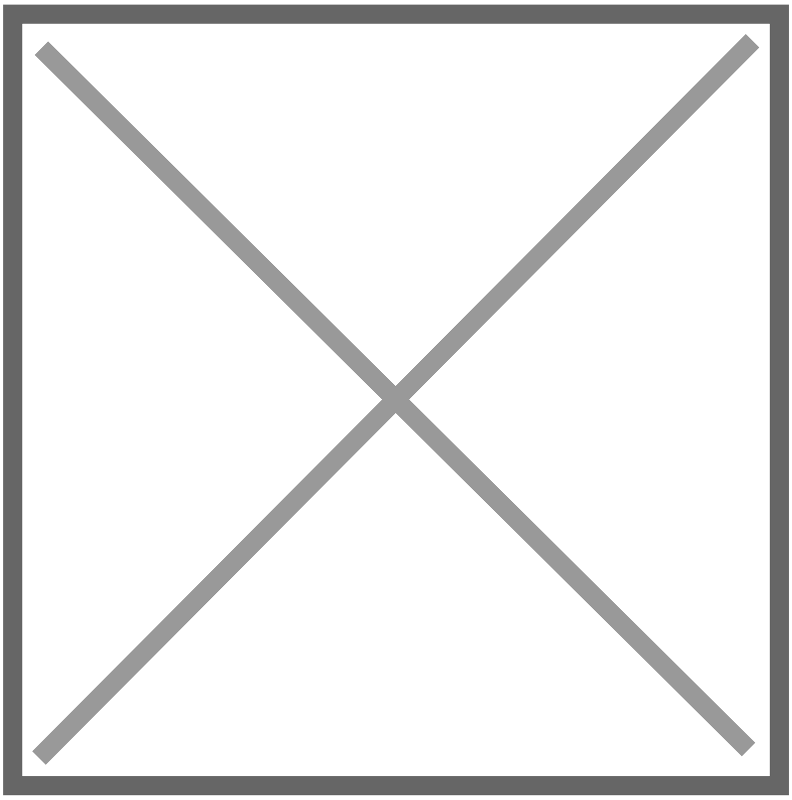

[illegible] H [illegible] loss [illegible] OBQ [illegible] OBS [illegible] 0 [illegible] 1bit [illegible] 0bit [illegible] OBQ [illegible]

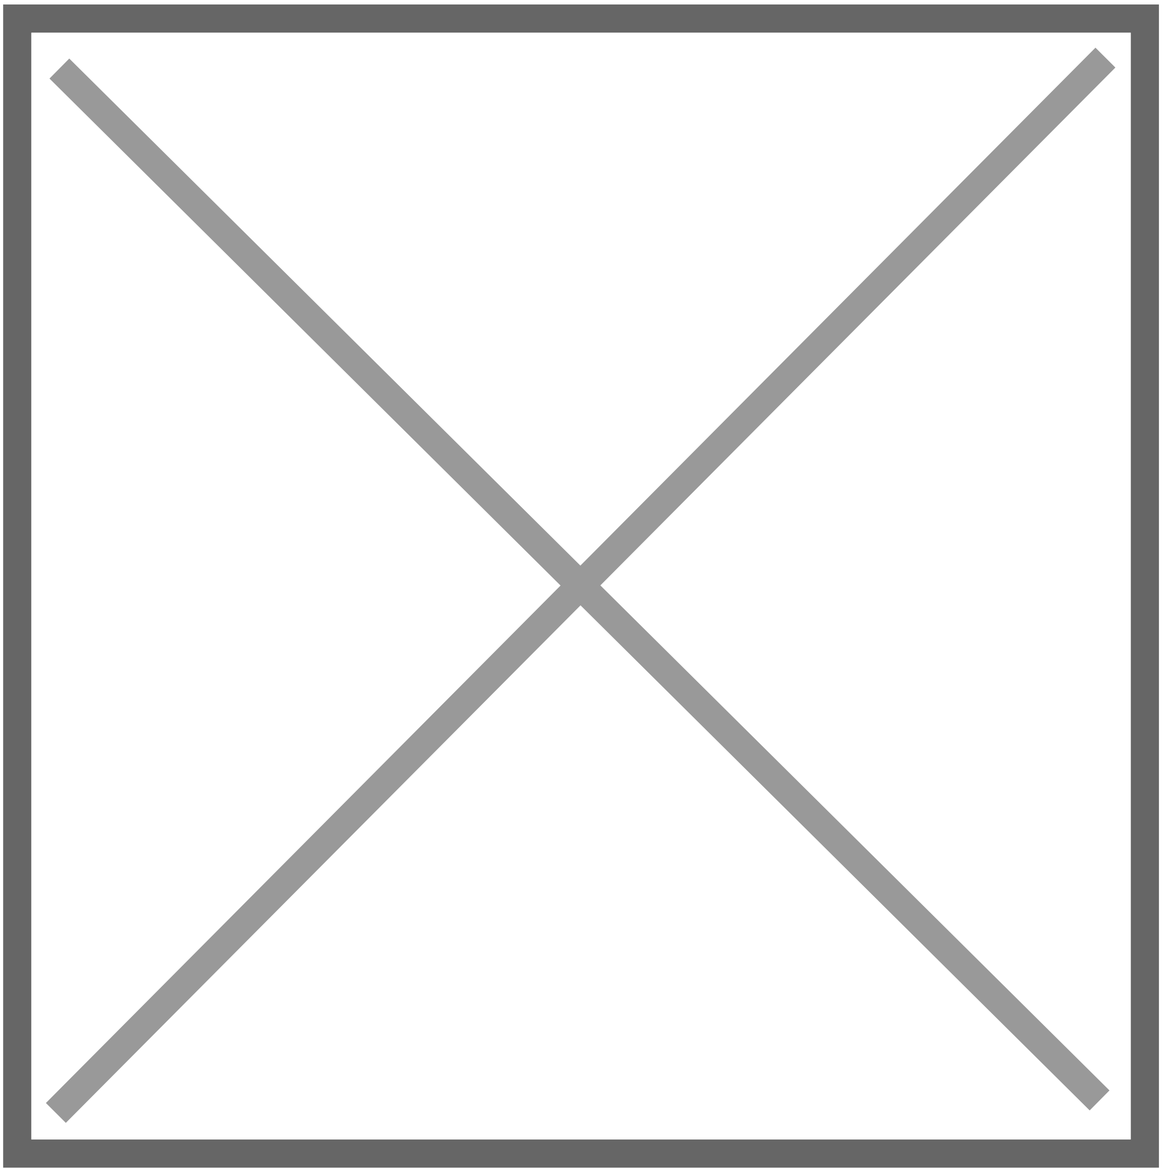

[illegible] quant(q) [illegible] q [illegible] quant grid [illegible]

quant(q) [illegible] OBS [illegible] OBQ [illegible]

[illegible] OBS

[illegible] q [illegible] / [illegible]

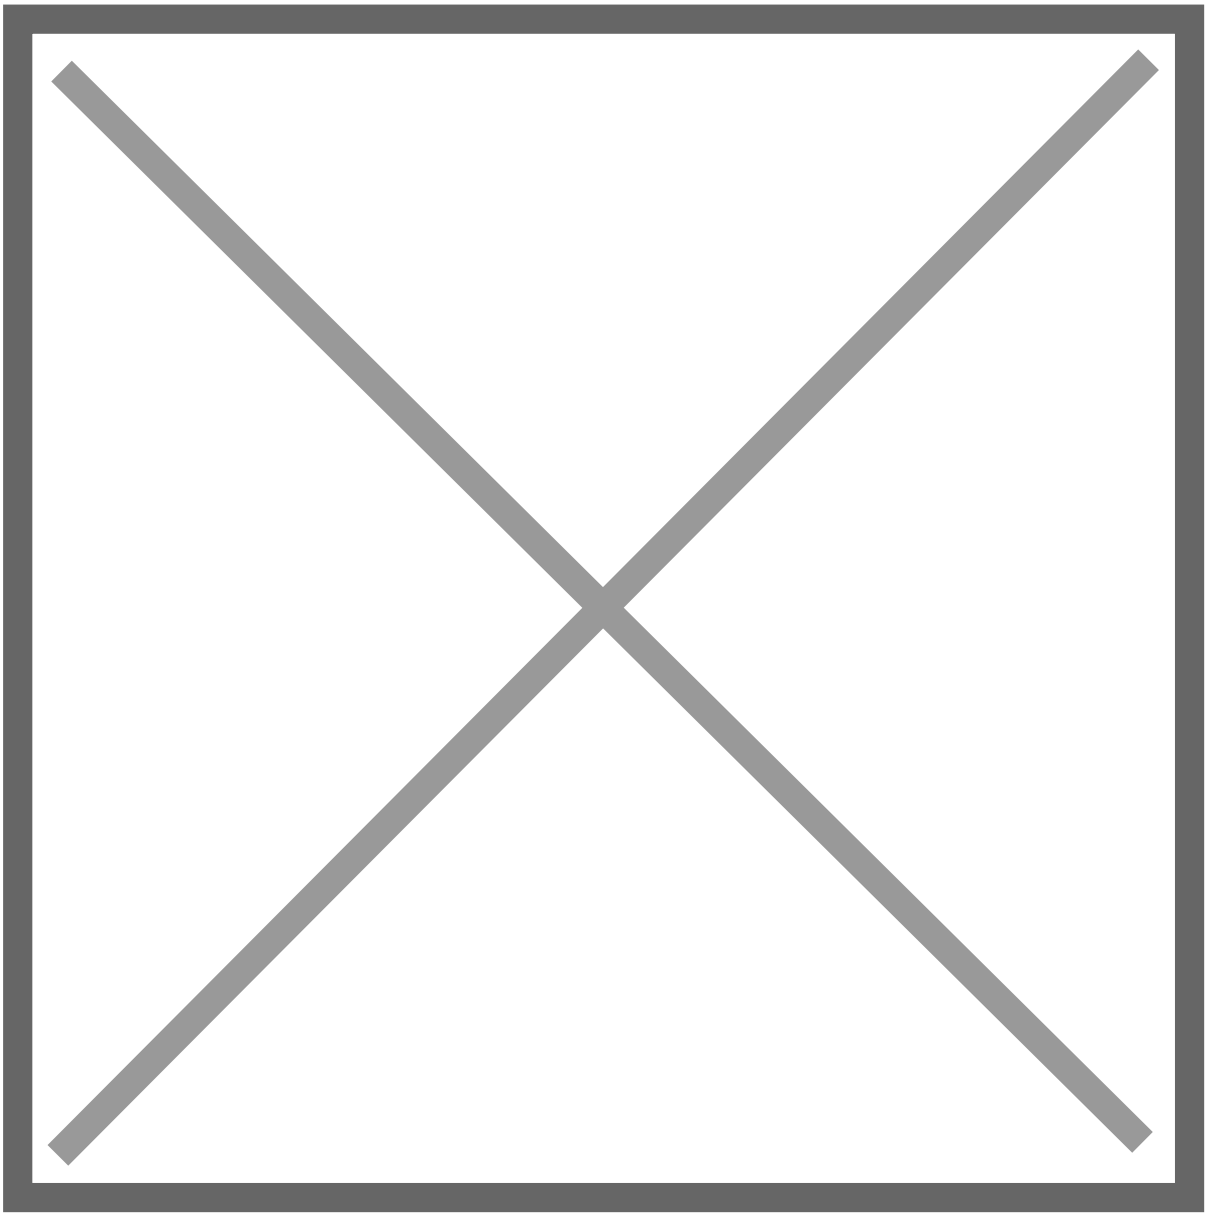

GPTQ OBS

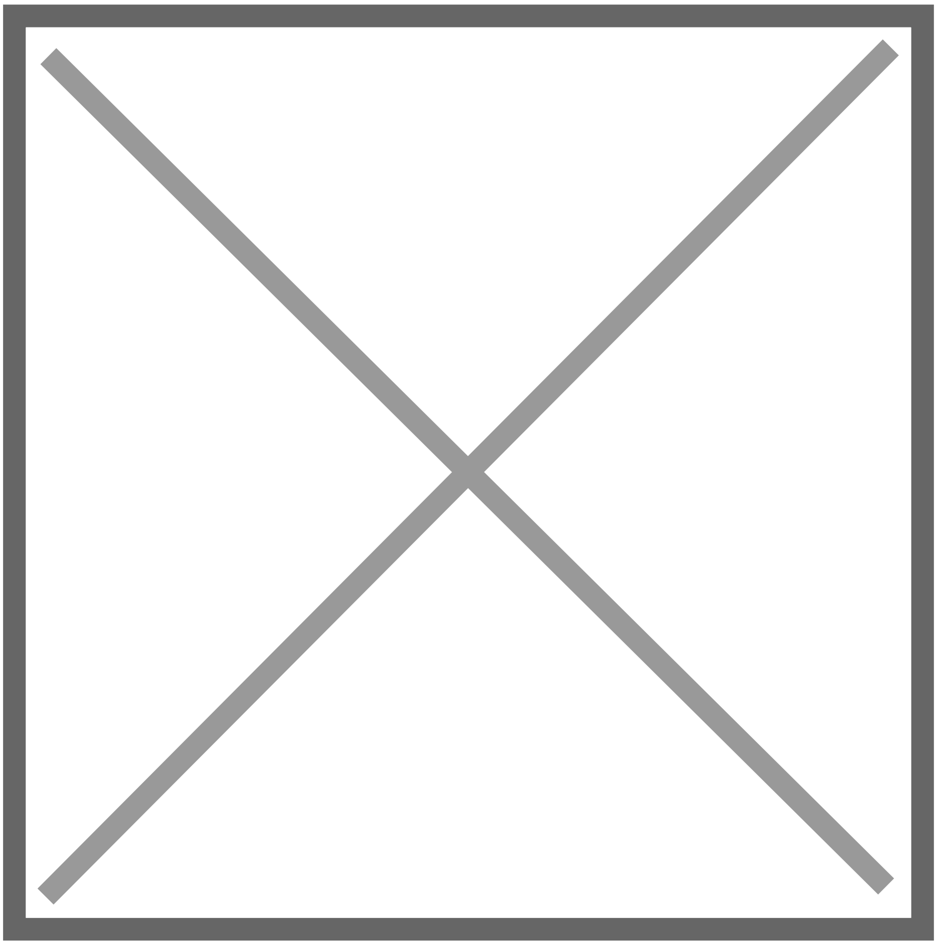

-
[illegible]
[illegible]

-[illegible]

-[illegible]
GPU[illegible]

-[illegible]

-[illegible] Scale[illegible] Zero[illegible] Group size
[illegible]

[illegible] GPTQ[illegible] LLM
[illegible]

SmoothQuant: W8A8[illegible] scale

[illegible] LLM[illegible] channel[illegible] activation outlier[illegible] (e.g. GLM-130B[illegible]
30%[illegible] int8[illegible] -128[illegible] 127[illegible] outlier[illegible] -128[illegible] 127

[illegible] 8bit [illegible] outlier [illegible] [-20,20] [illegible] [-10,10] [illegible] 8bit [illegible] 5bit [illegible] 4bit [illegible]

[illegible]

[illegible] INT8 [illegible] INT4 [illegible] **weight** [illegible] **activation** [illegible] **activation** [illegible] **weight** [illegible]

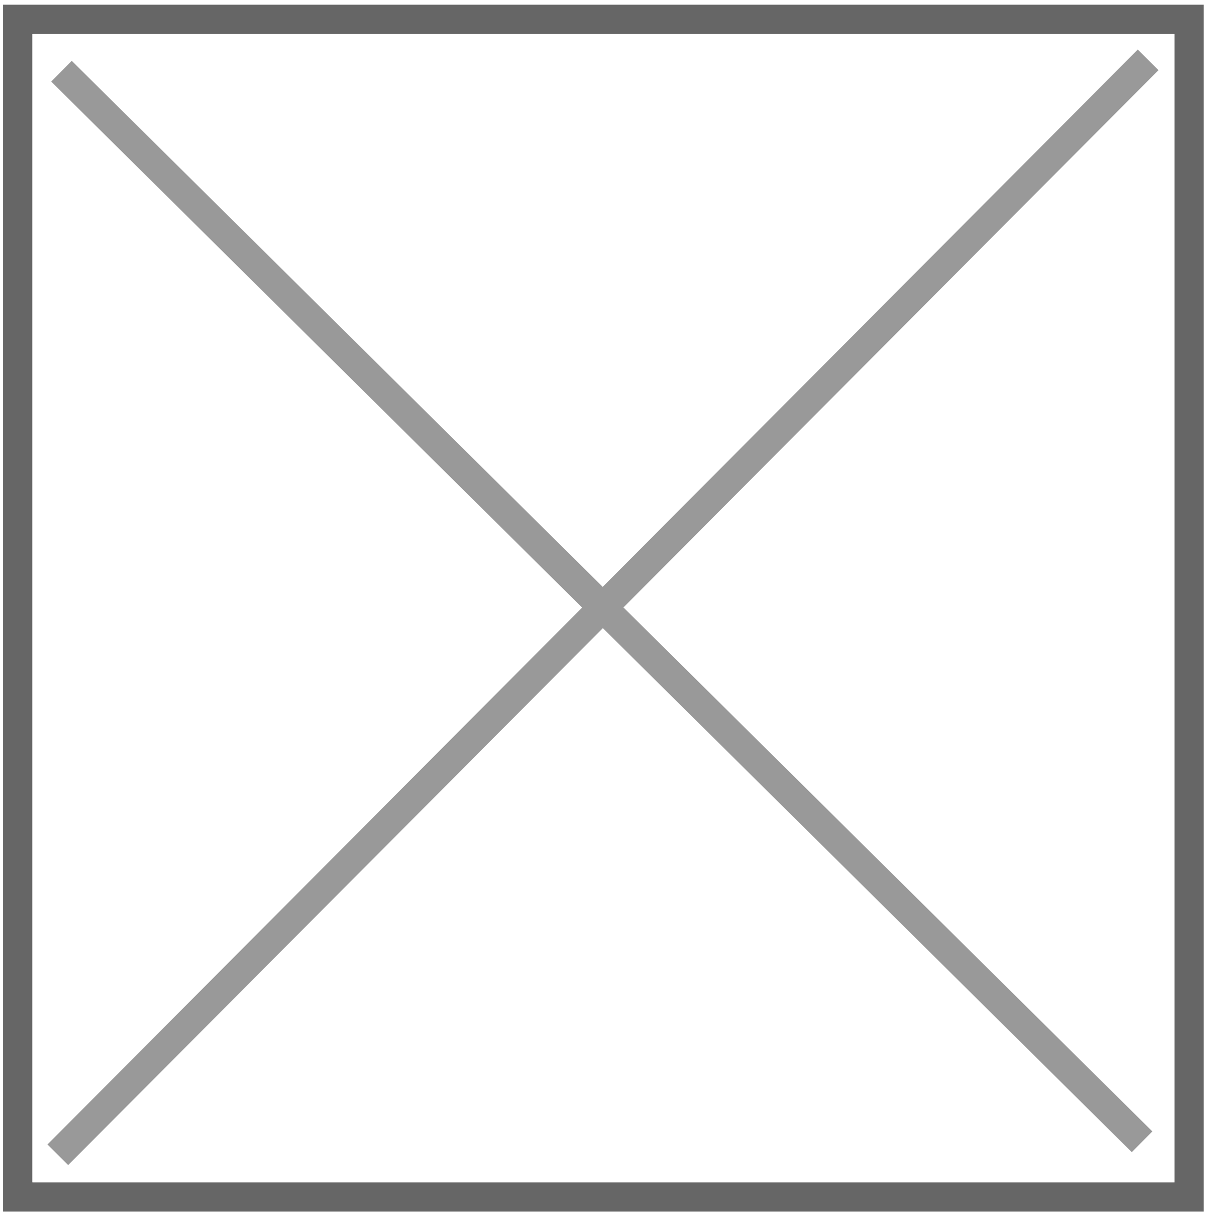

SQ [illegible] activation [illegible] weight [illegible] fp32 [illegible] X [illegible] act [illegible] W [illegible] weight [illegible] s [illegible] s [illegible]

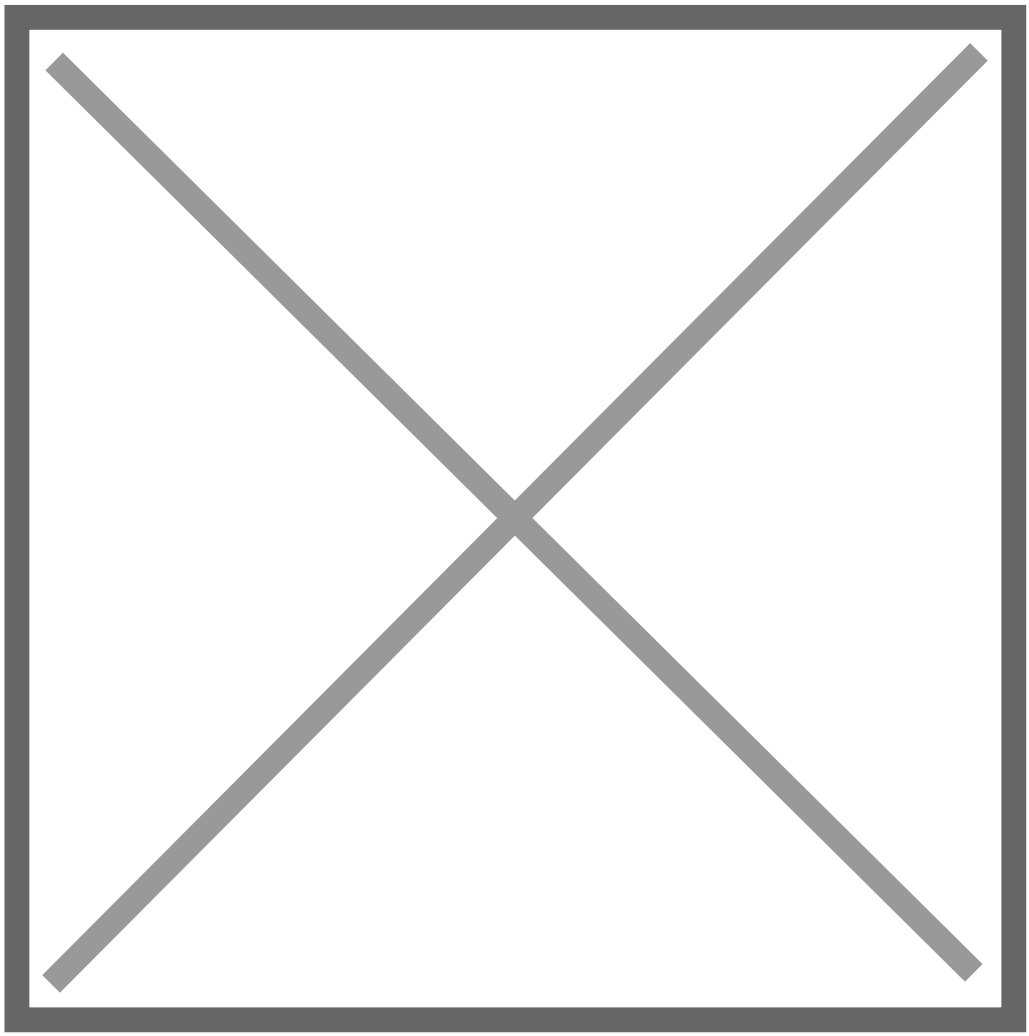

□□□□□□□□□□□□□

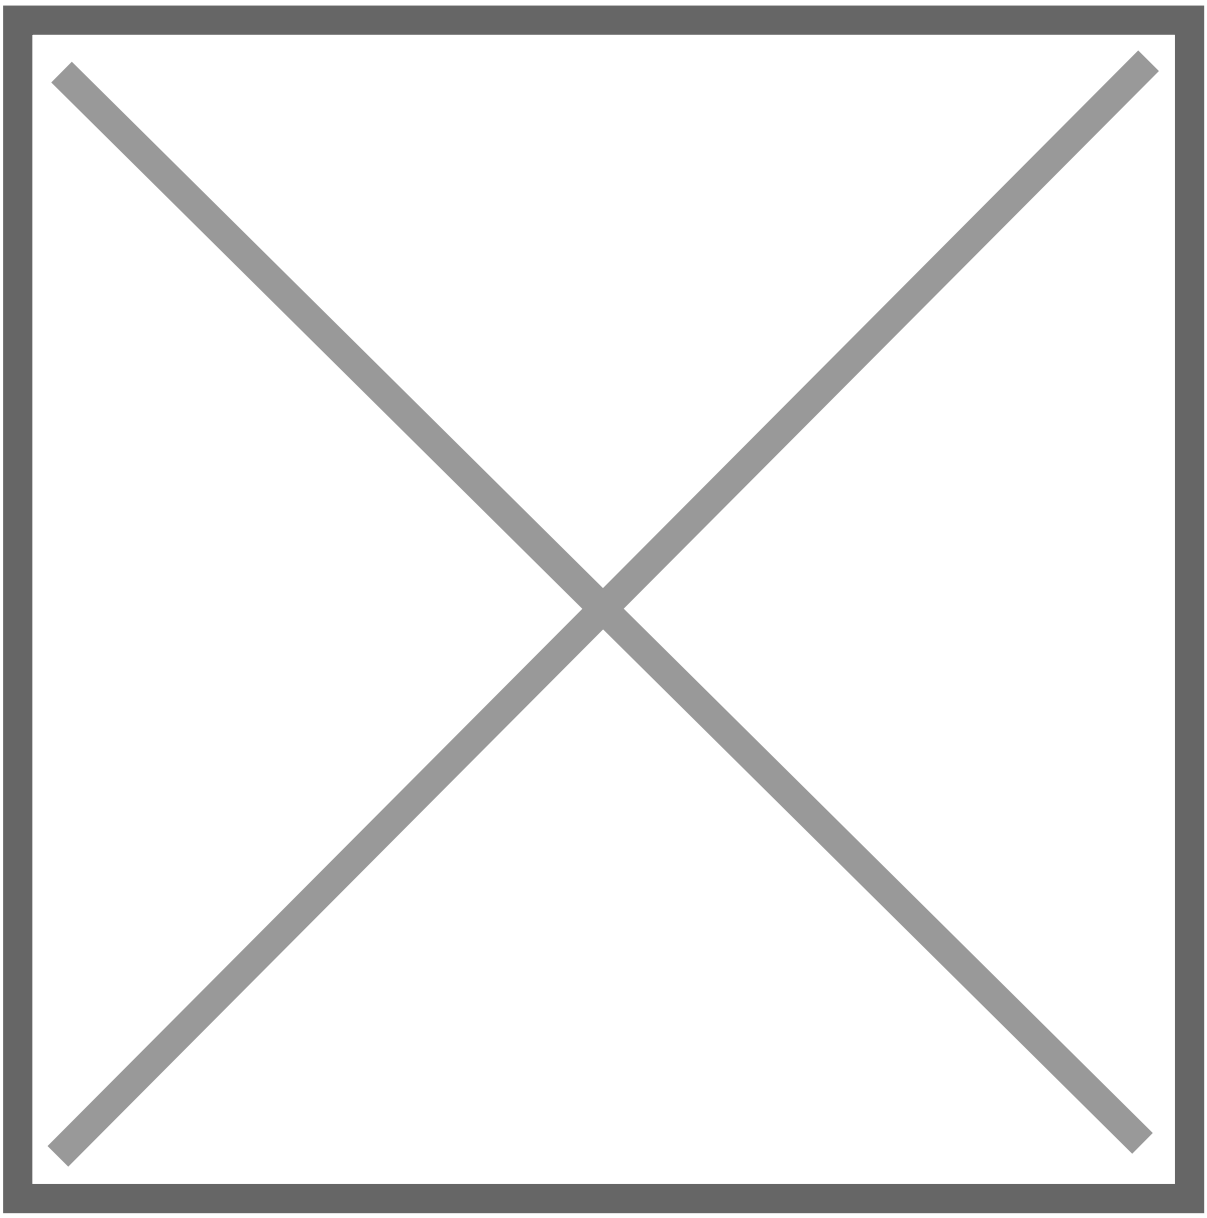

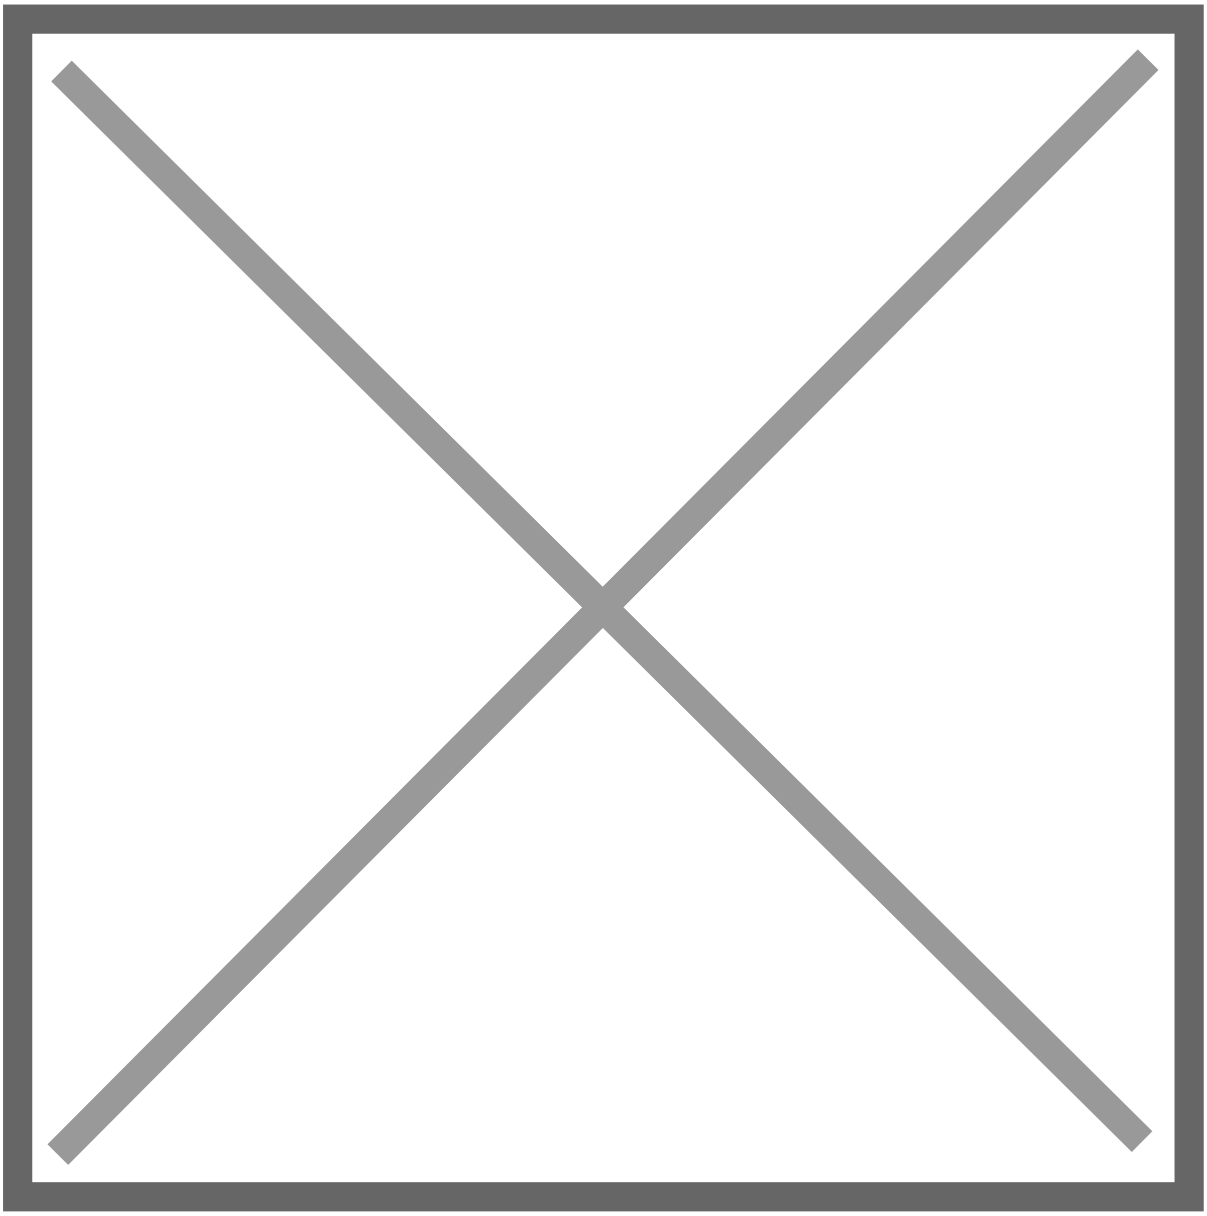

[illegible] LLM[illegible] weight-only[illegible] SmoothQuant[illegible] LLM [illegible] SmoothQuant[illegible] W8A8 [illegible] INT4[illegible] bit [illegible]

[illegible] SmoothQuant[illegible] **LLM** [illegible] weight-only [illegible] SmoothQuant[illegible] TensorRT-LLM[illegible] TensorRT[illegible] LLM [illegible] GPT2[illegible] SmoothQuant[illegible]

AWQ: [illegible]

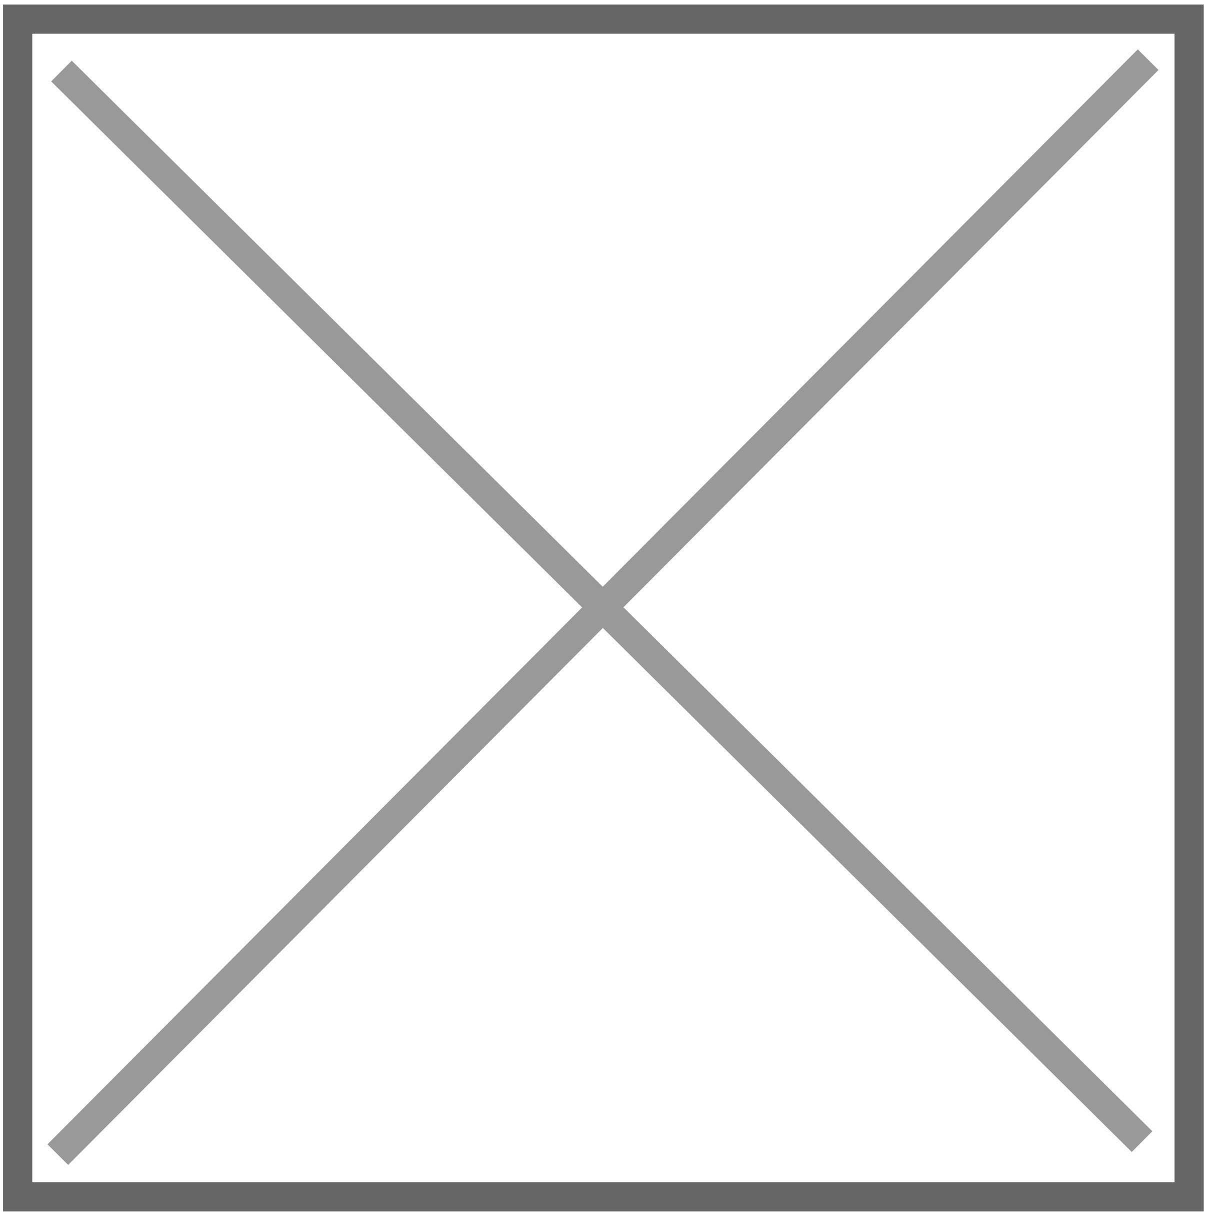

[illegible] [illegible] (a)[illegible] 1%[illegible] channel[illegible] fp16[illegible] (b)[illegible] (int8+fp16)[illegible] LLM.int8()[illegible] tensor [illegible] (c): [illegible] ***1%***[illegible] ***group*** [illegible] ***SmoothQuant*** [illegible]

[illegible] s[illegible] X [illegible]

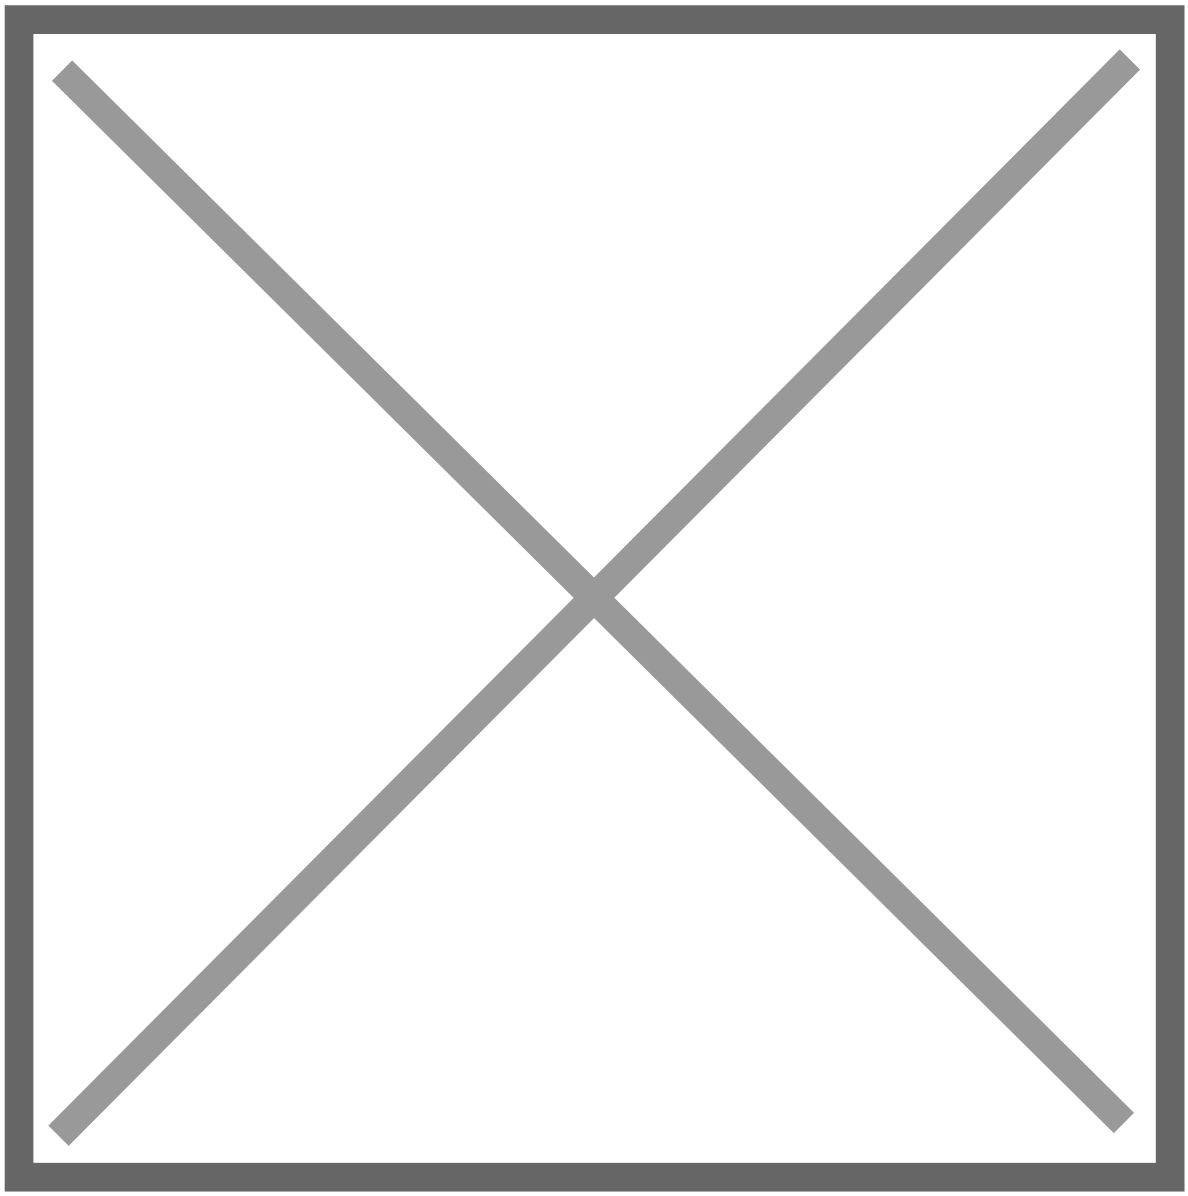

□□□□□□□□ s□□□□□□□□□□□□

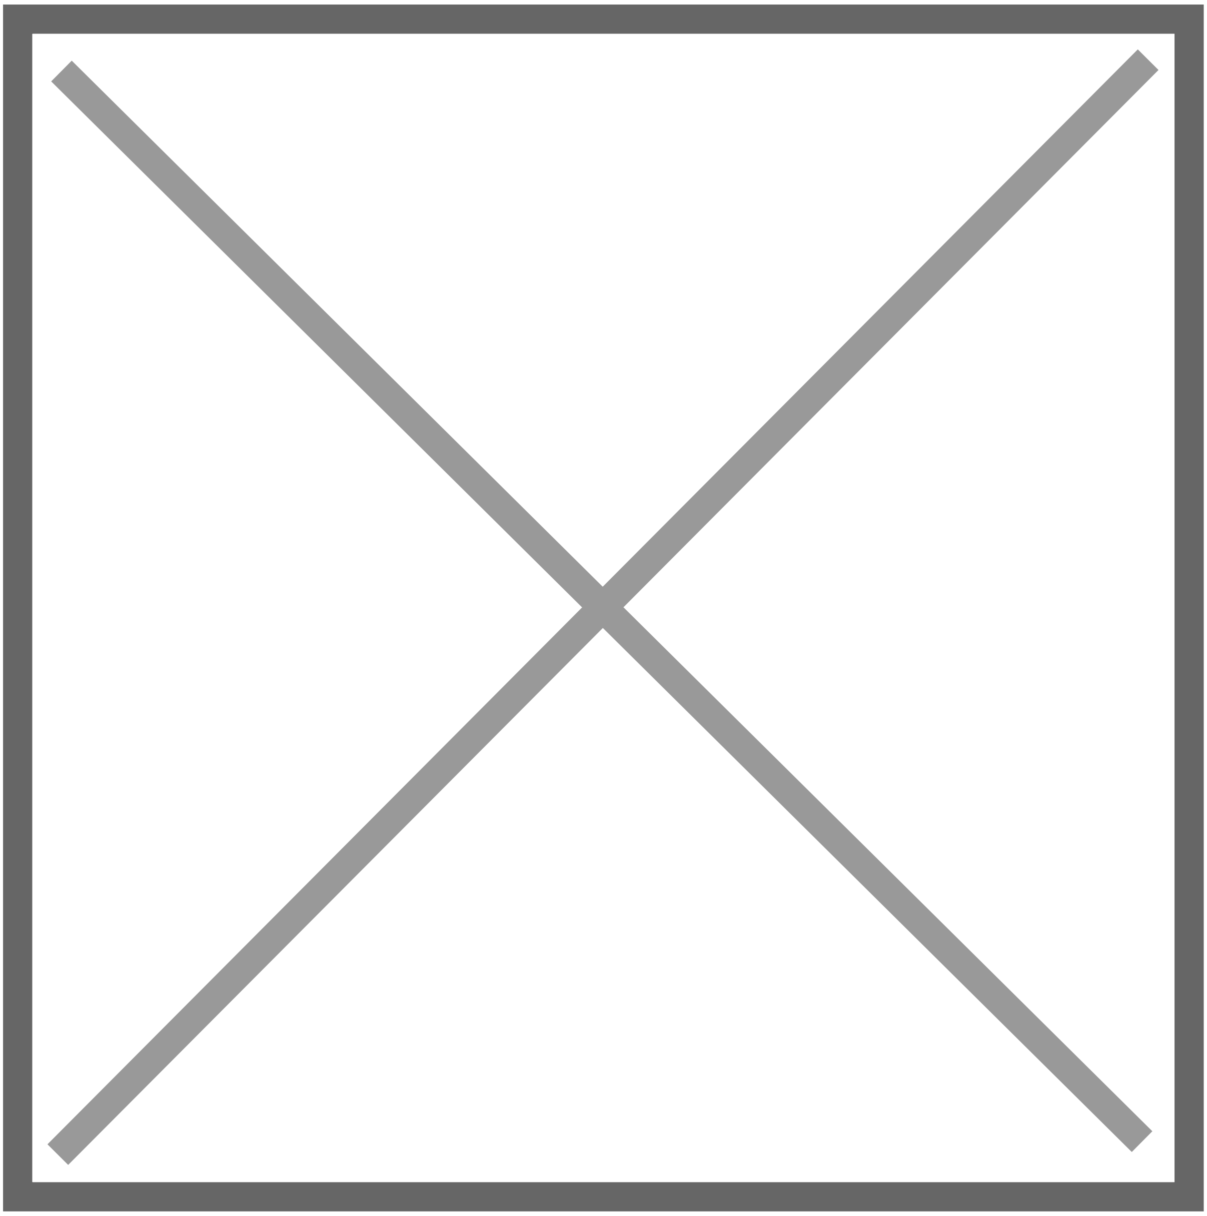

[illegible] sX = meanc_out|X|[illegible] sW = meanc_out|W|[illegible] sX[illegible] group [illegible] weight[illegible] group[illegible] α [illegible] β[illegible] [0, 1] [illegible] sW [illegible] group[illegible] weight [illegible]

[illegible] AWQ[illegible] kernel[illegible]

[illegible] smooth quant[illegible] AWQ[illegible] weight-only[illegible] SQ[illegible] W8A8[illegible] AWQ[illegible] GPTQ[illegible] kernel[illegible]

Llama. CPP[illegible] K-quant[illegible]

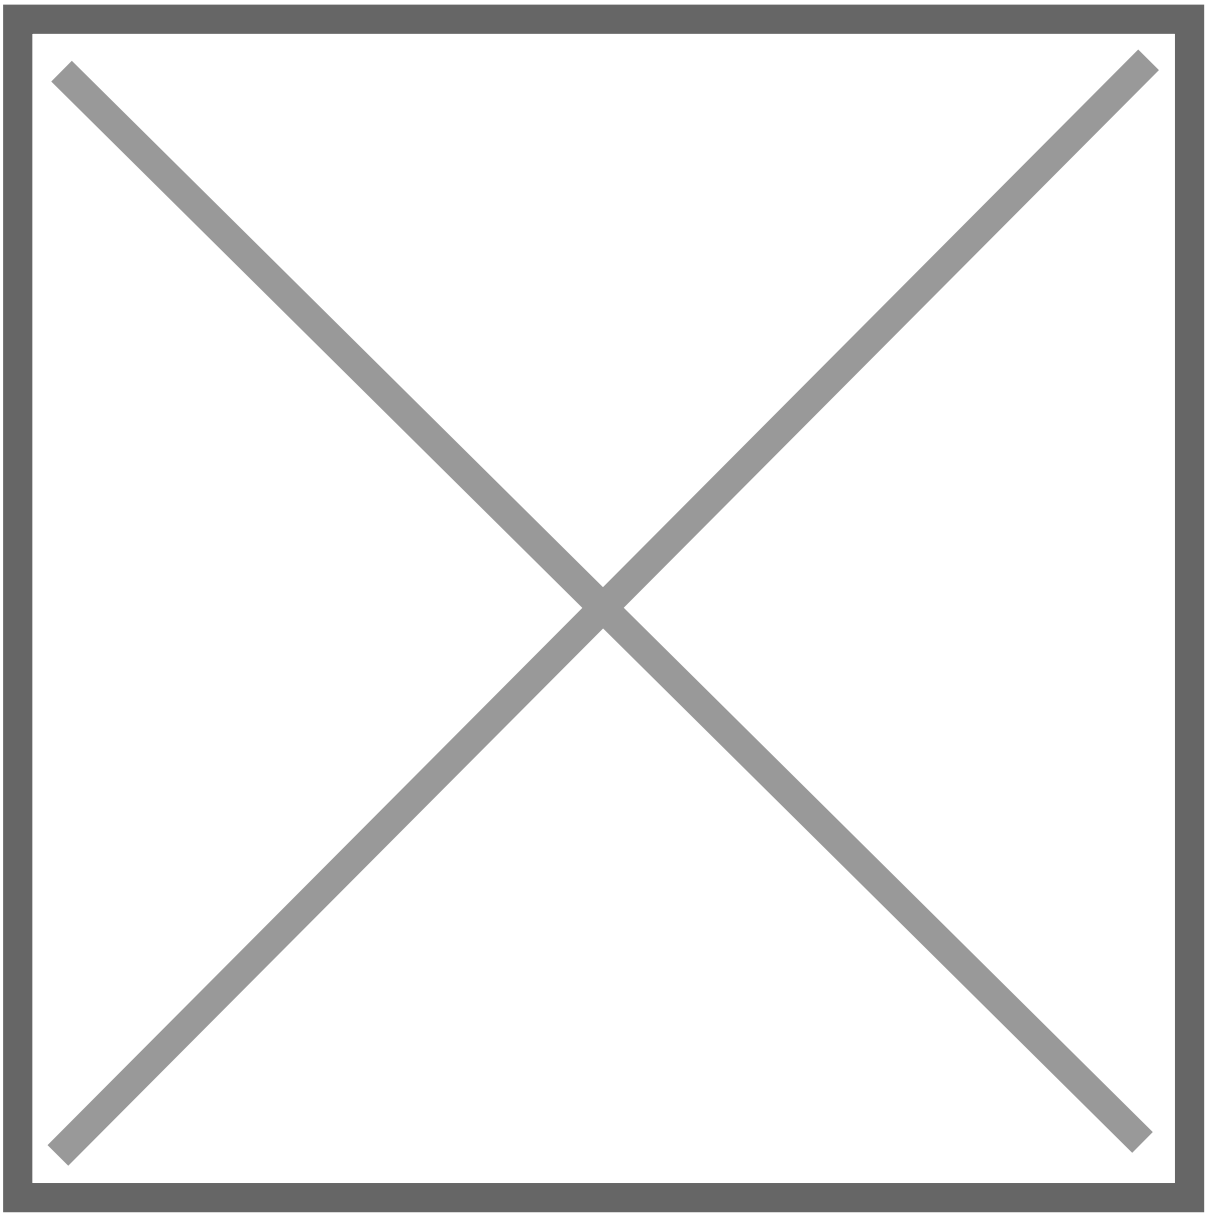

LLaMA.cpp项目是开发者Georgi Gerganov基于Meta释出的LLaMA模型实现的纯C/C++版本（基于其另一个项目：tensor库GGML），用于模型推理。llama.cpp的初衷是让开发者在没有GPU的条件下也能运行Meta的LLaMA模型。后来在他本人和社区开发者的支持下，也支持了cuBLAS的加速。在6月15日的PR中，llama.cpp正式支持了CUDA库的加速，用大量的cuda kernel加速了llama在GPU的推理速度。

llama.cpp中量化部分主要采用了两种方案：INT4量化和K-quant量化。量化选项如下图所示，其中Q后面的第一个数字x表示了量化到x bit，下划线后的第二个数字y为0时表示对称量化，没有零点。y为1时表示非对称量化，每个scale还有一个zero point，当y为K时表示K-quant量化方法，K后面的字母表示量化模型的参数规模：Small, Meduim，Large。例如，Q4_0和Q8_0选项代表进行INT4、INT8量化。

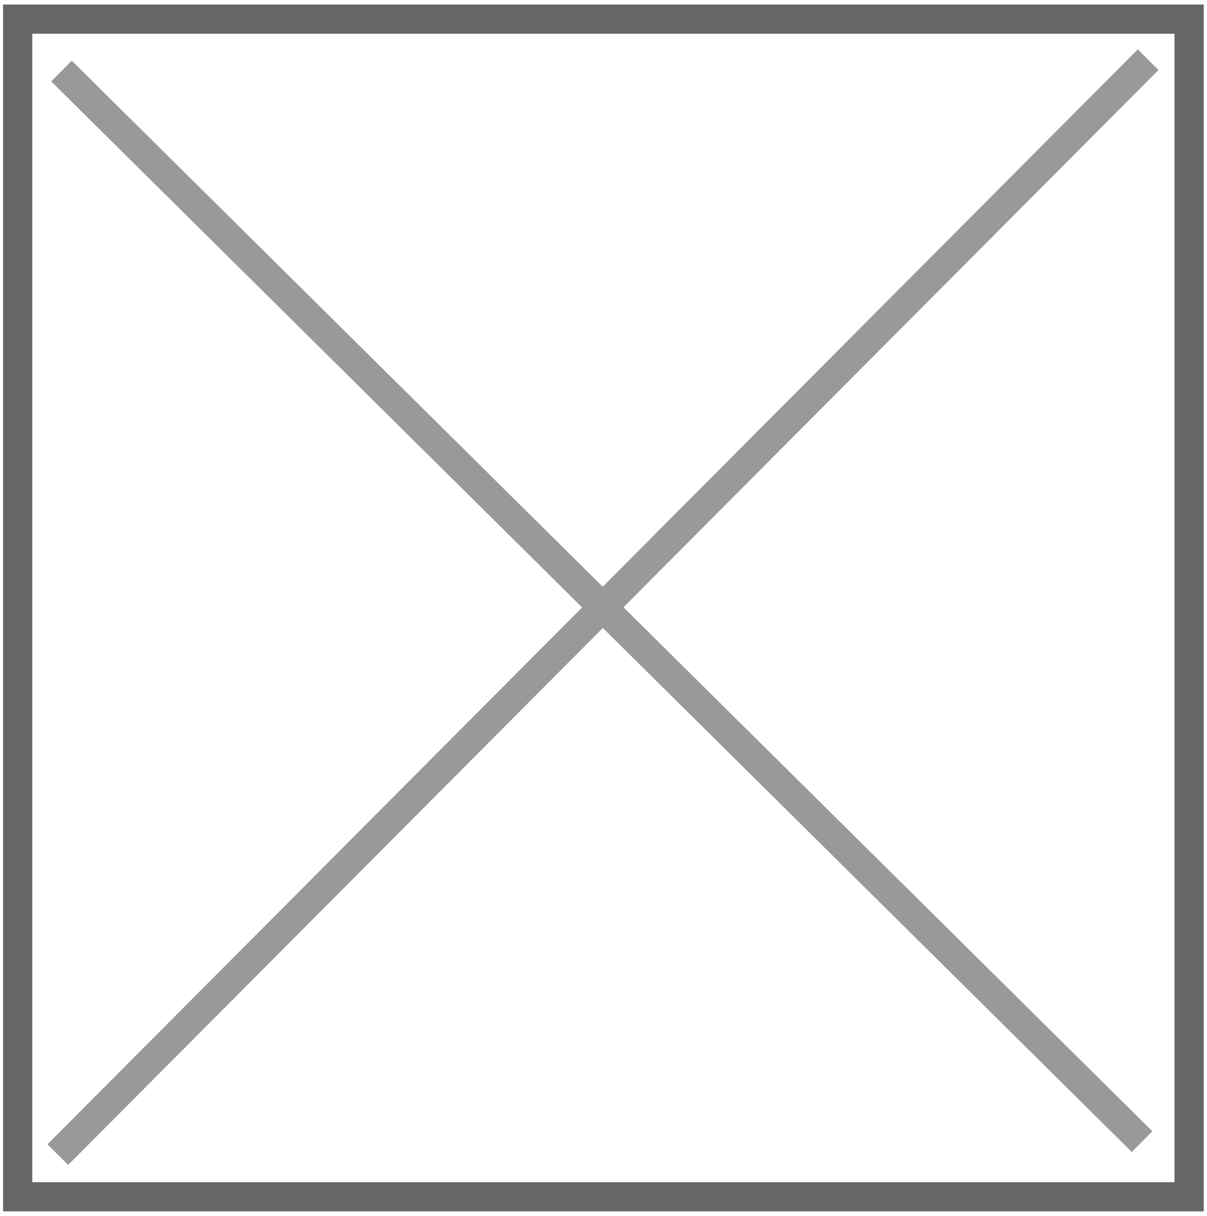

[illegible] y=K[illegible] K-quant[illegible] K-quant[illegible] 16 x 8[illegible] “[illegible]” [illegible] “[illegible]” [illegible] 16[illegible] 8 [illegible] scale[illegible] 16 [illegible] 1 [illegible] fp16 [illegible] K_2[illegible] 16 [illegible] **“[illegible]”**[illegible] size[illegible] 16*8[illegible] “[illegible]” K-quant[illegible] K_1[illegible] group size = 8, [illegible] 65B[illegible] 37.5GB [illegible]

[illegible] **llama.cpp**[illegible] **K-quant**[illegible] **fp16**[illegible] **3~4**[illegible] **(**[illegible] **token/s)**[illegible] fp16 [illegible] llama.cpp[illegible] c++[illegible] cuda [illegible] +gemm[illegible] GPT [illegible] KV Cache[illegible] token[illegible] vector-matrix[illegible] GEVM [illegible] matrix-matrix[illegible] GEMM[illegible] llama.cpp[illegible]

[illegible] llama.cpp[illegible] llama[illegible] **+K-quant**[illegible]

[illegible]

[illegible] llama.cpp [illegible] Kernel [illegible] llama.cpp [illegible]

QLora [illegible]

Lora [illegible] Qlora [illegible] QLora [illegible] lora [illegible]

[illegible]

[illegible] A [illegible] B [illegible] W [illegible] LoRA [illegible] AB [illegible] W [illegible] AB [illegible]

QLora [illegible]

Block-wise Quantization [illegible] block) [illegible] scale [illegible] block [illegible] 64, [illegible] K-quant [illegible]

Quantile Quantization [illegible] [illegible] Quantile Quantization. [illegible] Quantile [illegible]

[illegible] block-wise quantization [illegible] block [illegible] scale [illegible] float32, block [illegible] 64 [illegible] 32/64 = 0.5bit [illegible] 4bit [illegible] 0.5bit [illegible] 12.5% [illegible] scale [illegible] outliner [illegible] 256 [illegible] block [illegible] 8bit [illegible] c [illegible] 8/64 + 32/256 = 0.127 bit [illegible]

QLora [illegible] LLM [illegible]

[illegible]

- [illegible] LLM [illegible] weight-only [illegible]
- LLM [illegible] CNN [illegible] per-vector [illegible] + [illegible] scale [illegible]

- [illegible] fp16 [illegible] INT [illegible] Llama.cpp [illegible]
- LLM [illegible] INT4 [illegible] + [illegible] INT8 [illegible] INT4 [illegible]

Data type

data type

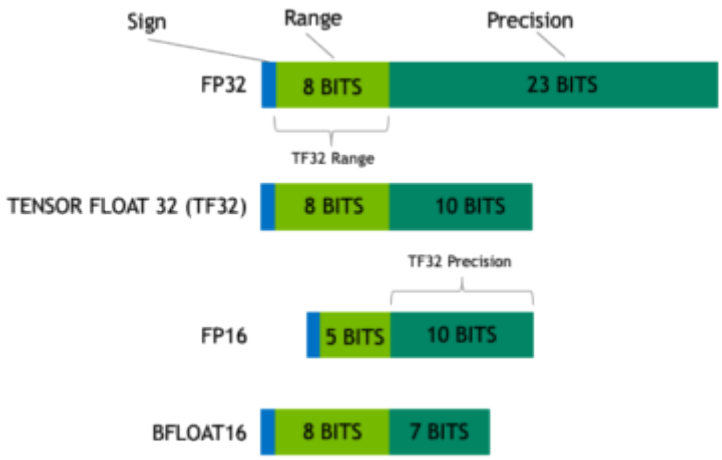
Sign
Range
Precision
FP32
8 BITS
23 BITS
TF32 Range
TENSOR FLOAT 32 (TF32)
8 BITS
10 BITS
TF32 Precision
FP16
5 BITS
10 BITS
BFLOAT16
8 BITS
7 BITS