
LLM-MQ: Mixed-precision Quantization for Efficient LLM Deployment

Shiyao Li^{1,2}, Xuefei Ning¹, Ke Hong^{1,2}, Tengxuan Liu^{1,2}, Luning Wang^{1,2},
Xiuhong Li^{2,3}, Kai Zhong¹, Guohao Dai^{2,4}, Huazhong Yang¹, Yu Wang¹

¹Dept. of EE, BNRist, Tsinghua University,

²Infinigence-AI,

³Peking University,

⁴Shanghai Jiao Tong University

Abstract

Large Language Models (LLMs) have demonstrated impressive performance across various tasks. Nevertheless, deploying LLMs on edge devices presents significant challenges, primarily due to their substantial model size (e.g., over 10 billion parameters). Low-precision quantization is a promising way to reduce the memory requirement of LLMs. However, directly applying ultra-low-bit quantization to LLMs leads to significant performance degradation and fails to meet a specific weight memory budget. In this paper, we propose LLM-MQ, a Mixed-precision Quantization method, to address the above issues. Our method mainly contains three folds: (1) We propose a **sparse outlier protection** strategy for low-precision layers by protecting the outliers in FP16 format to maintain the performance. (2) We propose **sensitivity-based precision allocation** to assign the proper bit-width for each layer within the given budget for weight memory based on their first-order information and quantization error. (3) We develop **efficient CUDA core kernels** to accelerate mix-precision LLMs by fusing the dequantization and General Matrix-Vector Multiplication (GEMV). With comparable performance on various tasks, LLM-MQ can flexibly quantize LLMs that meet the given budget for weight memory. On NVIDIA T4 GPU, we achieve up to $1.6\times$ end-to-end speedup compared to the pytorch FP16 baseline.

1 Introduction

Large language models (LLMs) have exhibited impressive performance across various language tasks. Recently, LLMs have led to the emergence of many interesting and helpful applications, such as ChatGPT (11), Copilot (2). However, deploying LLMs on edge devices has been challenging due to the enormous model size. For example, the LLaMA-2-70b model (15) needs at least 140 GB of memory for deployment. Even the NVIDIA A100 GPU with 80 GB memory is not enough to deploy this model. It is especially not enough for devices with limited memory capacities.

To meet the demands of edge inference scenarios where model size constraints are paramount, the memory overhead of weights is the main bottleneck. Consequently, low-precision weight-only quantization can effectively reduce the memory requirement of LLMs. Furthermore, numerous methods (8; 6; 12; 14; 7) have been developed, all of which employ a consistent low-bit-width representation across all layers. They successfully quantize the FP16 LLMs into 4-bit or 3-bit LLMs with acceptable performance. However, for edge devices with more limited memory budgets, it is challenging to push the bit-width for each layer to 2-bit without significant accuracy loss.

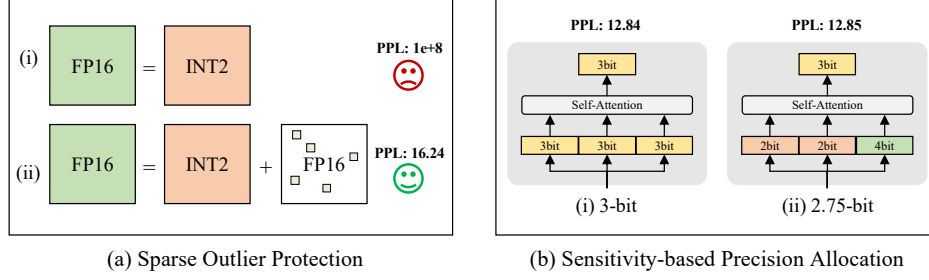


Figure 1: (a) Sparse Outlier Protection. (i) Baseline methods directly quantize the original FP16 weights into 2-bit, with dramatically accuracy loss (ii) We keep 0.5% outlier at FP16 and quantize the rest normal weights into 2-bit. (b) Sensitivity-based Precision Allocation. (i) Baseline methods use the same bit-width for each layer. (ii) We assign lower bit-width to the less sensitive layers.

In this paper, we delve deeper into mixed-precision quantization to address the above issue. Given that different layers exhibit varying sensitivities to quantization, assigning the appropriate bit-width to each layer can provide advantages in meeting the specified memory budget while maintaining minimal accuracy loss. The main contributions of this paper are:

- We propose a sparse outlier protection strategy to only quantize the normal weight into ultra-low precision and keep the outliers in FP16 format. Experiments show that we can significantly reduce the accuracy loss using 2-bit quantization.
- We propose a sensitivity-based precision allocation method by modeling the precision allocation as an integer programming problem. Specifically, we model the given budget for weight memory as a constraint of the proposed integer programming problem. Experiments show that we can push the average bit-width into 2.8-bit.
- We design efficient CUDA kernels to accelerate the decoding stage by fusing the dequantization and GEMV operations and also design a memory arrangement strategy for odd-bit weights to improve the speed of memory access further. On NVIDIA T4 GPU, we achieve up to $1.6\times$ end-to-end speedup compared to the pytorch FP16 baseline.

2 Method

2.1 Sparse Outlier Protection

As shown in Fig. 1 (a), while nearly 99.9% of the weights are concentrated around zero, the largest value is $\sim 10^7$ larger than the mean value. 2-bit uniform quantization results in a quantization error $2.3\times$ larger than 3-bit uniform quantization, leading to significant accuracy loss. To counteract this, we introduce sparse outlier protection, retaining 0.5% of outliers in FP16 while applying 2-bit quantization to the remaining values. This method decreases the error by 1.56 times, approximating 3-bit quantization outcomes. For efficient computation, FP16 outliers are stored using the CSR format, leveraging the sparse library (1), akin to the SqueezeLLM approach (7)."

2.2 Sensitivity-based Precision Allocation

The primary goal of quantization is to express model weights using low-bit-width representation, ensuring the changes in the model output remain minimal (5). Since different layers have different sensitivities, we propose to assign a high bit-width to high-sensitivity layers and a low bit-width to low-sensitivity layers in order to minimize the change in model output. To pinpoint the more sensitive layers, we employ the first-order Taylor approximation in Eq. 1 to determine how the model output changes in response to weight perturbations.

$$\mathcal{L}(Q_b(\mathbf{W}_i)) \approx \mathcal{L}(\mathbf{W}) + \mathbf{g}_i^T (\mathbf{W}_i - Q_b(\mathbf{W}_i)), \quad (1)$$

where \mathcal{L} is the loss function, \mathbf{g}_i is the gradient of the loss function with respect to the weight of i -th layer \mathbf{W}_i , and $Q_b(\cdot)$ is the b -bit quantization function. Note that in LLMs, the hessian matrix of $\sim 15\%$ layers are not positive semi-definite, which means these layers have not converged to a local minimum. In this case, the second-order information of LLMs struggles to show the change of loss function accurately. We focus on the first-order information:

$$\min |\mathcal{L}(Q_b(\mathbf{W}_i))| \leq |\mathcal{L}(\mathbf{W}_i)| + \min |\mathbf{g}_i^T (\mathbf{W}_i - Q_b(\mathbf{W}_i))| = |\mathcal{L}(\mathbf{W}_i)| + \min s_{i,b}, \quad (2)$$




Figure 2: (a) Directly apply 2-bit quantization to the original weight. (b) Only apply 2-bit quantization to the normal value while keeping 0.5% outliers in FP16 format.

where $s_{i,b}$ stand for $|\mathbf{g}_i^T(\mathbf{W}_i - Q_b(\mathbf{W}_i))|$. We minimize the sum of $s_{i,b}$ for all layers to minimize the change of loss function. In addition, we propose to use three different bit-width, that is 2, 3 and 4. We model the above bit-width allocation task as the following integer programming problem:

$$\arg \min_{c_{i,b}} \sum_i \sum_b^N c_{i,b} \cdot s_{i,b}, \quad (3)$$

$$\sum_b c_{i,b} = 1, \quad \sum_i \sum_b^N c_{i,b} \cdot \mathcal{M}(Q_b(\mathbf{W}_i)) \leq \mathcal{B}, \quad (4)$$

$$c_{i,b} \in \{0, 1\}, b \in \{2, 3, 4\}, \quad (5)$$

where Q_b is the b-bit quantization function, \mathcal{B} is the budget for weight memory of a certain device, \mathcal{M} is the function to calculate memory usage of weights, N is the layer number of the target LLM, and c_i is the one-hot indicator of i-th layer to determine its proper bit-width. We use the budget for weight memory as a constraint of the optimization objective. With the help of the efficient integer programming solver, we can find a proper bit-width allocation scheme within a few seconds.

2.3 Efficient GPU Kernel Design

To maximize the utilization of the GPU’s memory resources, we store two 4-bit weights in one byte. Eight 3-bit weights become three bytes, and four 2-bit weights are represented by one byte. To accelerate the inference of quantized models, we integrate weight dequantization with GEMV in our implementation, reducing memory access overhead. Additionally, we utilize the CUDA core instead of the Tensor core as employed in AWQ (8), primarily for two reasons: (1) The Tensor Core requires that each dimension of the inputs be at a minimum of eight, which is not required in the CUDA Core. In the GEMV operator, the minimal dimension of the inputs is one. Using zero-padding to expand the minimal dimension to 8 introduces inefficiencies. (2) The latency of GEMV latency is primarily dictated by model weight memory access, and the enhanced computational capacity of the Tensor Core does not alleviate this primary bottleneck.

3 Experiments

We focus on weight-only quantization with a group size of 128 (i.e., every 128 data points share a scaling factor and a zero point.). We also use the reparameterization strategy introduced in AWQ (8). We conduct experiments of the proposed LLM-MQ on various benchmarks using LLaMA2 (15) and OPT (17) families. We assess the performance of quantized models across several benchmarks. The evaluation code is based on lm-harness-evaluation¹. These include five zero-shot benchmarks: PIQA (3), HellaSwag (16), WinoGrande (13), ARC-e (4), and OpenBookQA (10). Additionally, we evaluate the model on a language modeling task using the Wiki benchmark (9). Tab. 1 presents the primary results for LLaMA2-13B, while the performance of other LLMs are detailed in the Appendix. Furthermore, in Tab. 1, we have not factored in the FP16 sparse weights when calculating the average bit-width, as they have a minimal impact, increasing it by a maximum of 0.08 bits, which is negligible.

The effect of the sparse outlier protection. As shown in Tab. 1, in the 3-bit and 4-bit quantization, RTN (Round-To-Nearest), AWQ and SqueezeLLM have no significant performance degradation.

¹<https://github.com/EleutherAI/lm-evaluation-harness>

Table 1: The performance of the LLaMA2-13B model on five zero-shot benchmarks by reporting their average accuracy and one language modeling task using perplexity.

Method	#Bit	PIQA	Hella.	Wino.	Arc-e	OpenBookQA	Avg. (\uparrow)	Wiki (\downarrow)
LLaMA2-13B	16.0	79.11	76.59	69.77	57.91	42.00	65.08	7.89
RTN	4.0	79.16	75.90	69.61	56.69	42.00	64.67	8.12
	3.0	77.15	73.68	67.64	56.90	40.80	63.23	9.26
	2.0	56.47	37.94	51.30	30.77	32.20	41.74	1056.33
AWQ (8)	4.0	79.00	76.14	70.32	57.49	42.20	65.03	8.08
	3.0	77.91	74.62	69.77	56.02	41.80	64.02	8.81
	2.0	50.76	25.69	50.28	26.98	31.60	37.06	5e6
SqueezeLLM (7)	4.0	78.94	76.05	69.14	57.70	42.80	64.93	8.44
	3.0	78.62	74.53	67.40	56.73	40.20	63.50	9.29
LLM-MQ (Ours)	4.0	79.49	76.31	69.30	58.50	41.20	64.96	8.03
	3.8	79.22	76.22	69.85	58.29	41.40	65.00	8.08
	3.6	79.05	75.88	69.77	58.59	42.20	65.10	8.23
	3.4	79.49	74.77	69.61	58.12	40.60	64.52	8.61
	3.2	79.33	75.12	67.96	57.87	41.60	64.38	8.43
	3.0	79.00	75.08	68.59	57.79	41.00	64.29	8.54
	2.8	78.73	74.32	67.96	57.95	41.20	64.03	8.83
	2.6	78.35	73.81	68.03	57.32	39.40	63.38	9.35
	2.4	77.31	72.93	68.59	54.63	40.00	62.69	10.03
	2.2	76.77	70.83	67.09	55.26	38.40	61.67	10.80
	2.0	75.84	68.32	65.51	54.29	37.20	60.23	12.17

However, in 2-bit quantization, the performance of RTN degrades significantly on all tasks. The performance of AWQ is even worse than RTN, especially on the Wiki benchmark. In LLM-MQ, we only protect the top 0.5% of the largest value in each layer in FP16 format and quantize the rest of the normal value to 2-bit. The average accuracy of LLM-MQ on zero-shot tasks is **23.17%** and **19.16%** better than AWQ and RTN. On the Wiki benchmark, the perplexity of LLM-MQ is more than **two orders of magnitude lower than** that of other methods.

The effect of the sensitivity-based precision allocation. As depicted in Tab. 1, for 3-bit and 4-bit quantization, RTN (Round-To-Nearest), AWQ and SqueezeLLM do not have significant performance degradation. As the average bit-width decreases, the performance across all tasks decreases gradually. Furthermore, the performance loss of LLMs is more significant at lower average bit-widths. We observed that when the average bit-width remains above 2.8 bits, the performance decline is notably gradual, with sporadic instances of slight improvement. While quantizing the model to the average bit-width of 3.6bit and 2.8bit, performance degradation remains under 0.5% in comparison to AWQ’s 4bit and 3bit models. In this case, if the target device is limited to deploying models with an average bit-width of 2.8 bits, the performance of the LLM-MQ 2.8-bit model significantly surpasses the 2-bit quantized models of AWQ and RTN.

The efficiency evaluation of the proposed CUDA kernels. In our implementation, the 2-bit and 3-bit kernels do not show significant speedup compared to the 4-bit kernel. The reason is that for every bit-width, we ensure each thread processes a fixed amount of data. Consequently, during 2-bit and 3-bit quantization, the utilization of GPU bandwidth is less efficient compared to 4-bit quantization. This hinders a higher acceleration ratio. We deploy our INT4 quantized LLaMA2-7B model on NVIDIA T4 GPU with 16 GB memory for evaluation. Compared to the original FP16 model, with the output token number of 32, 128, 512, 1024, and 2048, our quantized model with the proposed CUDA kernels demonstrate end-to-end acceleration factors of 1.27, 1.42, 1.65, 1.64, and 1.45, respectively. For the proposed GEMV Kernel, the 4096×4096 and 4096×11008 GEMV kernels are $2.33 \times$ and $2.48 \times$ faster than CUBLAS FP16 GEMV kernel.

Acknowledgement

This work was supported by the National Natural Science Foundation of China (No. U19B2019, 61832007, 62104128, U21B2031), Tsinghua EE Xilinx AI Research Fund, Tsinghua-Meituan Joint Institute for Digital Life, and Beijing National Research Center for Information Science and Technology (BNRist).

References

- [1] <https://developer.nvidia.com/cusparse>.
- [2] <https://github.com/features/copilot>.
- [3] Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439, 2020.
- [4] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- [5] Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Hawq-v2: Hessian aware trace-weighted quantization of neural networks. *Advances in neural information processing systems*, 33:18518–18529, 2020.
- [6] Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers, 2023.
- [7] Sehoon Kim, Coleman Hooper, Amir Gholami, Zhen Dong, Xiuyu Li, Sheng Shen, Michael W Mahoney, and Kurt Keutzer. Squeezellm: Dense-and-sparse quantization. *arXiv preprint arXiv:2306.07629*, 2023.
- [8] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Xingyu Dang, and Song Han. Awq: Activation-aware weight quantization for llm compression and acceleration, 2023.
- [9] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- [10] Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.
- [11] OpenAI. Gpt-4 technical report, 2023.
- [12] Gunho Park, Baeseong Park, Minsub Kim, Sungjae Lee, Jeonghoon Kim, Beomseok Kwon, Se Jung Kwon, Byeongwook Kim, Youngjoo Lee, and Dongsoo Lee. Lut-gemm: Quantized matrix multiplication based on luts for efficient inference in large-scale generative language models, 2023.
- [13] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- [14] Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Daniel Y. Fu, Zhiqiang Xie, Beidi Chen, Clark Barrett, Joseph E. Gonzalez, Percy Liang, Christopher Ré, Ion Stoica, and Ce Zhang. Flexgen: High-throughput generative inference of large language models with a single gpu, 2023.
- [15] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [16] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- [17] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.

Appendix

A. Experiment on more LLMs

Table 2: The performance of the LLaMA2-7B model on five zero-shot benchmarks by reporting their average accuracy and one language modeling task using perplexity.

Method	#Bit	PIQA	Hella.	Wino.	Arc-e	OpenBookQA	Avg. (\uparrow)	Wiki (\downarrow)
LLaMA2-7B	16.0	76.88	72.95	67.25	53.49	41.00	62.31	8.79
RTN	4.0	77.04	72.57	66.54	55.39	40.00	62.31	9.21
	3.0	75.52	71.10	67.01	52.78	40.20	61.32	11.21
	2.0	51.31	26.27	49.49	27.40	26.80	36.25	2e6
AWQ (8)	4.0	76.71	72.58	66.69	53.28	41.20	62.09	9.04
	3.0	76.66	70.66	65.43	52.65	40.60	61.20	10.30
	2.0	49.73	26.14	49.80	26.52	32.40	36.92	1e7
SqueezeLLM (7)	4.0	76.71	72.23	68.27	53.41	40.00	62.12	9.50
	3.0	76.50	69.95	65.90	51.47	40.40	60.84	10.56
LLM-MQ (Ours)	4.0	76.82	72.42	67.96	53.79	41.40	62.48	8.96
	3.8	77.09	72.35	67.01	53.41	41.20	62.21	9.04
	3.6	76.61	71.85	66.85	52.74	39.80	61.57	9.23
	3.4	75.95	71.55	67.25	51.56	41.00	61.46	9.34
	3.2	76.39	71.28	64.80	52.57	39.40	60.89	9.52
	3.0	76.28	71.03	65.98	51.47	40.00	60.95	9.65
	2.8	76.01	70.23	66.22	51.30	39.60	60.67	10.17
	2.6	75.03	69.53	64.96	50.13	40.40	60.01	10.96
	2.4	74.21	67.67	64.56	48.36	38.20	58.60	12.01
	2.2	74.48	65.51	63.85	47.43	37.60	57.77	13.37
	2.0	72.31	61.14	59.91	45.88	38.80	55.61	15.93

Table 3: The performance of the OPT-6.7B model on five zero-shot benchmarks by reporting their average accuracy and one language modeling task using perplexity.

Method	#Bit	PIQA	Hella.	Wino.	Arc-e	OpenBookQA	Avg. (\uparrow)	Wiki (\downarrow)
OPT-6.7B	16.0	76.44	67.17	65.19	60.10	37.20	61.22	12.29
RTN	4.0	76.33	65.69	64.09	58.59	37.60	60.46	13.02
	3.0	72.42	58.41	59.91	52.40	34.60	55.55	43.15
	2.0	49.46	26.08	47.43	24.92	27.60	35.10	2e5
AWQ (8)	4.0	76.39	66.84	64.56	60.10	36.60	60.90	12.44
	3.0	75.79	65.45	64.80	57.00	38.00	60.21	12.99
	2.0	71.06	56.72	59.91	52.74	35.00	55.09	18.77
SqueezeLLM (7)	4.0	75.79	65.86	63.54	60.27	37.20	60.92	12.45
	3.0	75.68	63.96	64.88	58.67	35.40	59.72	13.17
GPTQ (6)	4.0	75.79	65.86	63.54	59.64	38.00	60.57	12.88
	3.0	74.32	60.07	62.90	53.91	35.60	57.36	17.87
LLM-MQ (Ours)	4.0	76.39	67.16	65.11	60.19	37.80	61.33	12.41
	3.8	76.39	66.91	65.11	60.10	38.20	61.34	12.46
	3.6	76.39	66.32	64.40	59.60	38.20	60.98	12.97
	3.4	75.84	65.77	64.88	58.67	39.00	60.83	13.17
	3.2	75.73	65.34	65.43	58.92	36.60	60.40	13.48
	3.0	76.12	65.93	63.77	59.81	38.20	60.77	12.84
	2.8	75.84	65.71	63.30	59.60	38.60	60.61	13.32
	2.6	74.97	63.13	64.40	57.66	36.40	59.31	14.94
	2.4	75.52	64.14	64.96	58.71	36.40	59.95	14.64
	2.2	74.43	62.70	63.38	57.74	36.40	58.93	15.70
	2.0	74.27	60.87	61.88	55.93	35.20	57.63	17.09

Table 4: The performance of the OPT-13B model on five zero-shot benchmarks by reporting their average accuracy and one language modeling task using perplexity.

Method	#Bit	PIQA	Hella.	Wino.	Arc-e	OpenBookQA	Avg. (\uparrow)	Wiki (\downarrow)
OPT-13B	16.0	76.88	69.81	65.04	61.78	39.00	62.50	11.49
RTN	4.0	76.22	68.21	64.64	62.46	37.80	61.87	11.88
	3.0	70.51	45.57	58.17	50.17	33.20	51.52	45.36
	2.0	49.46	26.20	49.49	25.25	27.60	35.60	1e6
AWQ (8)	4.0	76.01	69.66	65.43	61.74	39.00	62.37	11.60
	3.0	76.55	68.25	64.56	59.97	36.60	61.19	12.03
	2.0	71.98	57.07	60.54	50.17	35.20	54.99	16.06
SqueezeLLM (7)	4.0	76.61	68.91	64.88	62.29	39.00	62.34	11.62
	3.0	75.30	66.36	64.72	59.01	38.40	60.80	13.37
GPTQ (6)	4.0	76.33	69.02	64.40	60.27	38.00	61.60	11.74
	3.0	74.86	64.46	62.43	56.06	35.80	58.72	13.07
LLM-MQ (Ours)	4.0	76.66	69.37	65.43	61.36	38.00	62.16	11.59
	3.8	76.77	69.39	65.27	60.94	39.00	62.27	11.62
	3.6	76.22	68.98	64.48	61.74	39.60	62.20	11.82
	3.4	76.71	68.86	64.88	62.04	39.20	62.34	11.92
	3.2	75.95	67.24	67.09	60.27	37.80	61.67	12.70
	3.0	76.22	68.43	64.64	62.08	39.40	62.15	11.99
	2.8	74.92	67.13	64.80	61.45	38.00	61.26	12.88
	2.6	75.35	66.88	64.56	60.02	36.60	60.68	13.39
	2.4	74.70	65.32	64.17	59.09	36.20	59.90	14.45
	2.2	74.59	62.59	65.35	57.58	35.00	59.02	15.79
	2.0	73.23	60.63	64.17	54.55	34.20	57.36	17.24