



HAL
open science

When Quantization Affects Confidence of Large Language Models?

Guillaume Metzler, Irina Proskurina, Julien Velcin, Luc Brun

► **To cite this version:**

Guillaume Metzler, Irina Proskurina, Julien Velcin, Luc Brun. When Quantization Affects Confidence of Large Language Models?. 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Jun 2024, Mexico City, Mexico. hal-04646224

HAL Id: hal-04646224

<https://hal.science/hal-04646224>

Submitted on 12 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

When Quantization Affects Confidence of Large Language Models?

Irina Proskurina, Luc Brun, Guillaume Metzler, Julien Velcin

Université de Lyon, Lyon 2, ERIC UR 3083, France

Correspondence: Irina.Proskurina@univ-lyon2.fr

Abstract

Recent studies introduced effective compression techniques for Large Language Models (LLMs) via post-training quantization or low-bit weight representation. Although quantized weights offer storage efficiency and allow for faster inference, existing works have indicated that quantization might compromise performance and exacerbate biases in LLMs. This study investigates the confidence and calibration of quantized models, considering factors such as language model type and scale as contributors to quantization loss. Firstly, we reveal that quantization with GPTQ to 4-bit results in a decrease in confidence regarding true labels, with varying impacts observed among different language models. Secondly, we observe fluctuations in the impact on confidence across different scales. Finally, we propose an explanation for quantization loss based on confidence levels, indicating that quantization disproportionately affects samples where the full model exhibited low confidence levels in the first place. We make our code and quantized models publicly available.¹

1 Introduction

Large language models (LLMs) are widely used in a variety of natural language generation applications (Bahdanau et al., 2014; Brown et al., 2020; Winata et al., 2021; Le Scao et al., 2022; Touvron et al., 2023). LLMs have been proven to achieve high performance in zero and few-shot prompting, providing results on par with fine-tuned baselines, especially in commonsense reasoning tasks (Zhang et al., 2022; Le Scao et al., 2022; Jiang et al., 2023). Kaplan et al., 2020 show that emerging abilities come with the scale increase, which makes well-performing larger models less accessible and limits their practical usability. A range of efficient compression and acceleration methods, including quan-

¹<https://github.com/upunaprosk/quantized-lm-confidence>

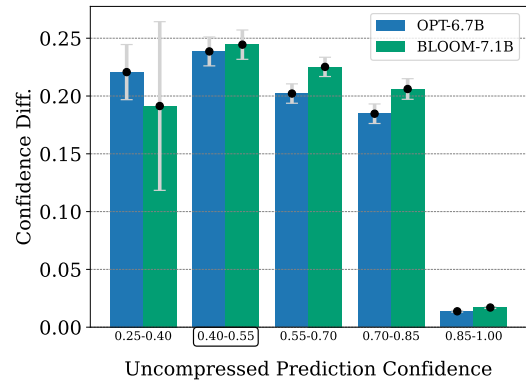


Figure 1: Quantization-induced absolute confidence shifts in original (pre-compression) low and high confidence samples (BLOOM and OPT models, HELLASWAG benchmark). The bin with the largest mean confidence shift is highlighted.

tization, have been developed that help to alleviate high latency and extensive storage demands (Gupta and Agrawal, 2020; Tao et al., 2022). Despite its efficacy as a compression technique, recent works show that quantization may degrade the initial performance and amplify the sensitivity of an LLM to certain linguistic phenomena and stereotypes (Liu et al., 2023; Ramesh et al., 2023). However, less attention has been paid to explaining the compression loss, particularly its variance across different texts. In this paper, we extend the existing research on the compression loss estimation; in particular, we measure the impact of quantization on the confidence of LLMs that can be initially overconfident in both right and wrong predictions (Jiang et al., 2021; Xiao et al., 2022b; Ahuja et al., 2022; Desai and Durrett, 2020).

Our main contributions are the following: (i) we investigate how quantization with GPTQ (Frantar et al., 2022) influences the calibration and confidence of LLMs, (ii) we assess the confidence alignment between compressed and full









Model	ARC EASY		BOOLQ		HELLASWAG		OPENBOOKQA		PIQA		XSTORY	
	Acc.↑	CE↓	Acc.↑	CE↓	Acc.↑	CE↓	Acc.↑	CE↓	Acc.↑	CE↓	Acc.↑	CE↓
 7B	81.10	7.94	83.61	38.62	61.30	34.3	32.60	45.24	80.83	45.24	78.89	4.78
 7B	75.25	9.99	75.05	38.78	56.94	37.8	34.0	44.56	78.67	44.94	76.77	4.97
 560M	47.35	29.13	55.14	26.91	31.58	64.81	17.2	61.16	64.09	40.98	61.22	5.13
1.1B	51.47	25.07	59.08	32.8	34.44	58.51	20.0	58.88	67.14	42.27	62.54	5.77
1.7B	56.31	21.99	61.77	38.29	37.54	55.67	21.40	56.64	68.77	41.4	64.66	5.65
3B	59.47	19.68	61.62	34.67	41.39	52.33	21.6	56.32	70.84	42.12	66.78	5.76
7.1B	65.03	15.57	62.81	32.28	46.49	48.54	25.20	53.23	72.63	42.52	70.55	5.53
 125M	43.56	32.76	55.44	30.13	29.18	62.84	16.6	61.19	62.00	41.51	58.84	5.9
350M	44.20	31.21	57.65	29.62	32.02	60.09	17.60	61.92	64.47	41.58	62.48	5.97
1.3B	56.99	20.85	57.67	26.42	41.56	52.7	23.4	55.04	71.71	41.49	70.28	5.6
2.7B	60.77	17.63	60.24	25.86	45.86	48.93	25.0	52.6	73.78	41.87	70.42	5.83
6.7B	65.57	15.58	66.05	28.05	50.51	45.25	27.6	50.99	76.28	43.72	73.6	5.62
13B	67.13	14.21	65.93	29.47	52.43	43.03	27.2	52.33	75.84	43.87	76.04	5.15

Table 1: Zero-shot accuracy scores (Acc.) and calibration error (CE) for full LLMs by benchmark with the difference in scores after quantization. We report expected CE for binary tasks and adaptive CE for multi-class benchmarks (ARC, BOOLQ, OPENBOOKQA). Notations: =MISTRAL; =LLAMA; =BLOOM; =OPT.

LLMs at scale, (iii) we explain the quantization loss from the initial confidence perspective.

Our null hypothesis is that the compressed vs. full predictive probability distributions are indistinguishable since prior work discussed a negligible accuracy drop in performance after quantization (Jacob et al., 2017; Dettmers et al., 2022; Xiao et al., 2022a). We analyze the relationship between models by comparing calibration scores—indicating a model’s ability to accurately reflect true probabilities—before and after quantization. To the best of our knowledge, our research is the first attempt to explain the quantization loss through the lens of predictive probabilities.

2 Related Work

The pretrained knowledge embedded in very large models has paved the way to parameter-efficient adaptation for downstream tasks, such as prompting and few-shot learning, bypassing the necessity for fine-tuning (Brown et al., 2020; Wei et al., 2022). The inference of LLMs can be accelerated through a low-bit representation of trained weights (*quantization*) and effective tensor slicing across multiple GPUs (DEEPSPEED (Rasley et al., 2020), ACCELERATE (Gugger et al., 2022), *inter alia*). Prior studies have estimated compression efficiency through: (1) latency-related measures determining throughput and a multiple of the original model’s inference speed-up, (2) the precision of weights approximation, and (3) performance decrease (gap) (Jacob et al., 2017; Dettmers et al., 2022; Xiao et al., 2022a; Frantar et al., 2022). Re-

cent comparative studies on interpreting compression loss have indicated that compression amplifies biases and stereotypes, highlighting a disparate quantization loss in multilingual LLMs across different architectures (Ramesh et al., 2023). In contrast, another line of research suggests that compression enhances fairness (Hessenthaler et al., 2022). Altogether, existing studies commonly measure compression loss by observing the deviation in performance before and after quantization. In this project, we adopt the recent GPTQ quantization method for compressing model weights and concentrate on the disparities between predictive probability distributions instead. For the first time, our approach reveals the relationship between the initial level of predictive confidence and quantization loss.

3 Methodology

We follow Jiang et al., 2020 and consider a classification problem where inputs to the model are questions x paired with candidate answers y to constitute concatenated sequences. The generative model then processes these concatenated question-answer pairs to predict the most probable answer \hat{y} from the provided choices Y for a given x :

$$\hat{y} = \arg \max_{y \in Y} p_{\text{LM}}(y|x).$$

Here, the probability of the token sequence y is derived as the product of individual token $y_{[i]}$ probabilities within the sequence, conditioned on x and

the preceding tokens $y_{[1:i-1]}$:

$$p_{\text{LM}}(y|x) = \prod_{i=1}^{|y|} p_{\text{LM}}(y_{[i]}|x, y_{[1:i-1]}),$$

where $|y|$ is the number of tokens composing the answer y .

For the entailment generation benchmarks, we use texts concatenated with possible completions as inputs to the model. We compare the quantized and full-precision models with the difference in the probabilities of the sequences $p_{\text{LM}}(y|x)$, further referred to as confidences.

3.1 Quantization

We quantize pre-trained weights of LLMs with a post-training quantization method known as GPTQ (OPTQ, Frantar et al., 2023). This approach employs iterative layer-wise weight quantization based on the input data, providing several benefits compared to other quantization methods: minimized weight approximation error, support for serialization across various bit configurations, and significantly accelerated inference using GPUs. We follow the GPTQ 4-bit configuration outlined by Frantar et al., 2023 and use a random sample of 128 sequences from the C4 dataset (Raffel et al., 2019) for quantization and set a grouping size equal to 128. Additional details regarding the quantization procedures can be found in Table 3 (Appendix A).

3.2 Evaluation

We focus on evaluating models’ confidence in predictions before and after quantization in a zero-shot setting. In an ideal scenario, we expect the model’s performance and confidence to remain consistent after quantization, preserving the initial calibration level. We evaluate the performance of LLMs post-compression using accuracy (Acc.) and calibration error (CE). For binary problems, we use the Expected Calibration Error (ECE; Naeini et al., 2015), calculated using reliability plots that bin predicted probabilities and compare them against actual accuracy. In multi-class benchmarks, we use the Adaptive Calibration Error (ACE; Nixon et al., 2019), which quantifies calibration performance by dividing predictions into equally sized bins based on confidence levels and comparing accuracy and confidence within these subsets.

Details regarding the binning parameters used are provided in Appendix B. We also examine two cases of miscalibration: (1) the model’s rejection

Model	Conf.	Conf _{err}	Conf _{true}	H
BLOOM	96.26	95.64	46.24	12.87
+ GPTQ	96.3	95.62	45.23*	12.89
OPT	96.51	95.57	50.37	12.12
+ GPTQ	96.5	95.55	49.78*	12.22
Mistral	96.85	95.02	61.14	10.96
+ GPTQ	96.89	95.13	59.73*	10.87
LLaMA	96.8	95.34	56.83	11.37
+ GPTQ	96.48	95.13	53.69*	12.21*

Table 2: Confidence and prediction entropy evaluation results on HELLASWAG for LLMs with $\sim 7\text{B}$ parameters. Quantized LLMs become less confident in both correct and wrong predictions. Conf.: Mean confidence in predictions; Conf_{err}: Mean confidence in wrong predictions; Conf_{true}: Mean confidence in true class; H=Mean predictive entropy in the answers, multiplied by 100. High entropy means that the model is more unsure about its predictions. The * denotes a significant difference with a confidence level set at 0.05 (paired t -test).

of correct predictions due to lower confidence and (2) the model’s incorrect prediction due to higher confidence. Specifically, we measure the model’s confidence Conf_{err} when predicting an incorrect class and the model’s confidence in the true class Conf_{true}.

4 Experiment Settings

Data We use six standard commonsense reasoning tasks for our analysis: ARC EASY (Clark et al., 2018), BOOLQ (Clark et al., 2019), PIQA (Bisk et al., 2020), HELLASWAG (Zellers et al., 2019), OBQA (OpenBookQA; Mihaylov et al., 2018), and XSTORY-EN (Mostafazadeh et al., 2017). These benchmarks vary in the types of language inference abilities assessed in LLMs: (1) question answering involving reading comprehension (BOOLQ), (2) natural text entailment (XSTORY-EN, HELLASWAG), (3) science fact knowledge (ARC, OBQA), and (4) physical commonsense (PIQA).

Models We use the following *causal* (auto-regressive) LLMs in our experiments: (1) BLOOM (Le Scao et al., 2022), (2) OPT (Zhang et al., 2022), (3) Mistral-7B (Jiang et al., 2023), and (4) LLaMA-7B (Touvron et al., 2023). To examine how confidence loss varies across different scales, we use various configurations of LLMs: BLOOM with 560M, 1.1B, 1.7B, 3B, and 7.1B parameters, and OPT with 125M, 350M, 1.3B, 2.7B, 6.7B, and 13B parameters.

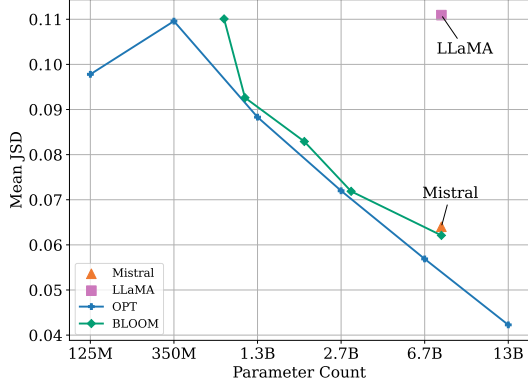


Figure 2: Mean Jensen-Shannon distances between full and quantized LLMs across benchmarks. The distances depict dissimilarities in true-class probability distributions.

5 Results

We conduct a series of experiments to estimate the impact of quantization on various aspects of LLMs’ performance, including calibration error, prediction entropy, cases of maximum confidence change, and the distribution dissimilarities between full and compressed models. We find variance in quantization impact across different families of models and their sizes, suggesting that scale and pre-training directly affect the further quantization loss.

Calibration Impact Table 1 outlines the classification results after quantization, evaluated through calibration error and accuracy metrics, along with the variation of these scores compared to the uncompressed LLMs. The general trend is that quantization amplifies the pre-existing high calibration error present in the models before compression across different models and benchmarks. This trend remains consistent across various model families, notably affecting the LLaMA-7B, which experiences a $\sim 10\%$ increase in pre-compression calibration error on the HELLASWAG dataset. Overall, scores associated with the HELLASWAG dataset are more significantly impacted compared to those of the BOOLQ and PIQA benchmarks.

Confidence Impact Table 2 presents the results obtained from four models, each having a near-equivalent number of parameters. Notably, across all models, a consistent trend of overconfidence emerges in both pre- and post-quantization stages, with an average confidence level around ~ 0.95 for incorrect predictions. Our analysis further shows

a statistically significant impact of quantization on the confidence associated with true-class predictions. Additionally, we observe an increase in entropy for the quantized LLMs shown in Table 5 (see Appendix D). This increase suggests an amplification in the variance across answers, reflecting increased uncertainty in answer selection due to quantization.

Identifying Cases of Confidence Change To identify instances of confidence change, we segment the models’ predictions into bins and calculate the confidence changes after quantization within each bin. In Figure 1, we depict the mean confidence changes for the BLOOM and OPT models on the HELLASWAG benchmark. The plot illustrates that samples with lower pre-quantization confidence levels are significantly affected by the quantization process, whereas samples in which the original model was confident show less impact. This observation suggests that quantization predominantly influences the confidence of samples where the original model exhibited lower confidence levels.

Jensen-Shannon Distances To illustrate the extent of differences between the distributions of the full and compressed models, we plot the mean Jensen-Shannon distances across benchmarks in Figure 2. These distances reflect the dissimilarity between the true-class probability distributions of the models. We find that the distances between original and compressed decrease as the model size scales up. Notably, most model families show a consistent trend in this behavior, except for LLaMa, which diverges from the patterns observed in other models of similar size ($\sim 7B$).

6 Conclusion

This paper investigates the impact of quantization on the confidence and calibration of LLMs. We demonstrate that quantization leads to an increase in calibration error and statistically significant changes in confidence levels for correct predictions. Through a detailed examination of confidence shifts, we identify instances of confidence change occurring in data where models lack confidence before quantization. Overall, our findings provide insights into quantization loss and suggest a potential direction for future work, emphasizing the need to focus on calibrating LLMs, specifically on uncertain examples. For example, future work

may focus on integrating the models' calibration, such as temperature scaling, into the quantization pipeline. Also, we have demonstrated that different model families, including LLAMA, MISTRAL, BLOOM, and OPT, exhibit varying degrees of susceptibility to quantization, as measured by changes in confidence levels. This suggests another direction for future research – benchmarking LLMs based on their response to quantization-induced confidence shifts.

Limitations

Our quantization techniques are currently limited to 4-bit post-training quantization with GPTQ. However, future work can benefit from exploring training-aware quantization approaches, studying different quantization factors, such as 2- and 3-bit weight representation, and quantization of activations.

In our evaluations, we employ zero-shot techniques, enabling the estimation of the pure quantization effect. Previous studies mentioned in related work included a fine-tuning step, whereas our approach avoids it. Yet, future work could involve few-shot analysis since this method has the potential to amplify or compensate for confidence and quantization loss.

Further research could apply our analysis to other generative tasks. Instead of predictive distributions over labels, one could consider those across tokens. This means using the full model's predictions as references and comparing the confidence in these generations after the quantization process.

Acknowledgments

This work has benefitted from access to the HPC resources provided by IDRIS under the allocation AD011014384, granted by GENCI, which facilitated the utilization of the Jean Zay supercomputer. Additionally, this research is supported by the ANR project Diké (No. ANR-21-CE23-0026-02).

References

Kabir Ahuja, Sunayana Sitaram, Sandipan Dandapat, and Monojit Choudhury. 2022. [On the calibration of massively multilingual language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4310–4323, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *CoRR*, abs/1409.0473.

Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. [PiQA: Reasoning about physical commonsense in natural language](#). In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.

Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? Try ARC, the AI2 reasoning challenge](#). *CoRR*, abs/1803.05457.

Shrey Desai and Greg Durrett. 2020. [Calibration of pre-trained transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.

Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. [LLM.int8\(\): 8-bit matrix multiplication for transformers at scale](#). *ArXiv*, abs/2208.07339.

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2022. [GPTQ: Accurate post-training quantization for generative pre-trained transformers](#). *ArXiv*, abs/2210.17323.

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2023. [OPTQ: Accurate quantization for generative pre-trained transformers](#). In *International Conference on Learning Representations*.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023. [A framework for few-shot language model evaluation](#).

Sylvain Gugger, Lysandre Debut, Thomas Wolf, Philipp Schmid, Zachary Mueller, Sourab Mangrulkar, Marc

- Sun, and Benjamin Bossan. 2022. Accelerate: Training and inference at scale made simple, efficient and adaptable. <https://github.com/huggingface/accelerate>.
- Manish Gupta and Puneet Agrawal. 2020. Compression of deep learning models for text: A survey. *ACM Trans. Knowl. Discov. Data*, 16:61:1–61:55.
- Marius Hessenthaler, Emma Strubell, Dirk Hovy, and Anne Lauscher. 2022. Bridging fairness and environmental sustainability in natural language processing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7817–7836, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew G. Howard, Hartwig Adam, and Dmitry Kalenichenko. 2017. Quantization and training of neural networks for efficient integer-arithmetic-only inference. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2704–2713.
- Albert Qiaochu Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L'elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7B. *ArXiv*, abs/2310.06825.
- Zhengbao Jiang, J. Araki, Haibo Ding, and Graham Neubig. 2020. How can we know when language models know? On the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeff Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *ArXiv*, abs/2001.08361.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. BLOOM: A 176b-parameter open-access multilingual language model. *ArXiv*.
- Peiyu Liu, Zikang Liu, Ze-Feng Gao, Dawei Gao, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji rong Wen. 2023. Do emergent abilities exist in quantized large language models: An empirical study. *ArXiv*, abs/2307.08072.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? A new dataset for open book question answering. In *Conference on Empirical Methods in Natural Language Processing*.
- Nasrin Mostafazadeh, Michael Roth, Annie Louis, Nathanael Chambers, and James Allen. 2017. LS-DSem 2017 shared task: The story cloze test. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pages 46–51, Valencia, Spain. Association for Computational Linguistics.
- Mahdi Pakdaman Naeni, Gregory F. Cooper, and Milos Hauskrecht. 2015. Obtaining well calibrated probabilities using bayesian binning. *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, 2015:2901–2907.
- Jeremy Nixon, Michael W. Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. 2019. Measuring calibration in deep learning. *ArXiv*, abs/1904.01685.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Krithika Ramesh, Arnav Chavan, Shrey Pandit, and Sunayana Sitaram. 2023. A comparative study on the impact of model compression techniques on fairness in language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15762–15782, Toronto, Canada. Association for Computational Linguistics.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Chaofan Tao, Lu Hou, Wei Zhang, Lifeng Shang, Xin Jiang, Qun Liu, Ping Luo, and Ngai Wong. 2022. Compression of generative pre-trained language models via quantization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4821–4836, Dublin, Ireland. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and efficient foundation language models. *ArXiv*, abs/2302.13971.

- Dennis Ulmer, Jes Frellsen, and Christian Hardmeier. 2022. [Exploring predictive uncertainty and calibration in NLP: A study on the impact of method & data scarcity](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2707–2735, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed Hui hsin Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *ArXiv*, abs/2206.07682.
- Genta Indra Winata, Andrea Madotto, Zhaojiang Lin, Rosanne Liu, Jason Yosinski, and Pascale Fung. 2021. [Language models are few-shot multilingual learners](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 1–15, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Julien Demouth, and Song Han. 2022a. [SmoothQuant: Accurate and efficient post-training quantization for large language models](#). *ArXiv*, abs/2211.10438.
- Yuxin Xiao, Paul Pu Liang, Umang Bhatt, Willie Neiswanger, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2022b. [Uncertainty quantification with pre-trained language models: A large-scale empirical analysis](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7273–7284, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona T. Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [OPT: Open pre-trained transformer language models](#). *ArXiv*, abs/2205.01068.

A Quantization Parameters

Parameter	Value
Num bits	4
Group size	128
Dampening factor (%)	0.01
Desc act	false
Symmetry	true
True sequential	true

Table 3: Configuration for GPTQ

B Evaluation Details

In this section, we provide further details on the used measures for the experiments.

Jensen-Shannon Divergence In Figure 2, we give the distance dissimilarities in the true-class probability distributions using the Jensen-Shannon divergence. For a given dataset, we focus on the true class probabilities, $p \in \mathbb{R}^n$, for the full model, and $q \in \mathbb{R}^n$ for the quantized one, where n denotes the number of instances.

The Jensen Shannon-Divergence between these two distributions is defined by:

$$\begin{aligned} JSD(p, q) &= \frac{1}{2} \left(KL \left(p \parallel \frac{p+q}{2} \right) + KL \left(q \parallel \frac{p+q}{2} \right) \right), \\ &= \sum_{i=1}^n p_i \ln \left(\frac{2p_i}{p_i+q_i} \right) + q_i \ln \left(\frac{2q_i}{p_i+q_i} \right), \end{aligned}$$

where KL denoted the Kullback-Leibler divergence and p_i and q_i are the true-class probabilities of the i -th instance for the full and quantized model respectively.

These distances are then averaged over all the studied datasets.

Expected Calibration Error (ECE) Let us consider a model h , which assigns confidence (which are probabilities) of belonging in a given class. These confidence scores can be divided into several bins $B_m, m = 1, \dots, M$ where M is the number of bins. More precisely, an instance belongs to the bin B_m if its confidence score in the true class $conf_i$ is in a given range (e.g. if

$(m-1)/M \leq conf_i \leq m/M$). In a given bin B_m , we can also measure the accuracy of the model, i.e., compute the ratio of instances in the bin B_m that are well-classified.

The *expected calibration error* is then defined as the weighted mean, where the weights depend on the number of instances in the bin of the absolute difference between the accuracy $acc(B_m)$ of the bin and the mean confidence score in the bin $\overline{conf}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} conf_i$, i.e.,

$$ECE = \sum_{m=1}^M \frac{|B_m|}{n} |acc(B_m) - \overline{conf}(B_m)|,$$

where n is the sample size. Note this error has been developed for binary classification tasks and can be extended to multi-class settings using the so-called *SCE* (Nixon et al., 2019), but this first extension has been shown to be not relevant for all studies (Ulmer et al., 2022). The authors rather use the *adaptive calibration error*, which works with equal size bins.

Adaptive Calibration Error (ACE) The *adaptive calibration error* is defined by

$$ACE = \frac{1}{CM} \sum_{c=1}^C \sum_{m=1}^M |acc(B_m, c) - \overline{conf}(B_m, c)|,$$

where C is the number of classes, M is the number of bins that are created, $acc(B_m, c)$ is the accuracy on class c in the m -th bin and $\overline{conf}(B_m, c)$ is the mean confidence score for class c in the m -th. In this case, all the bins have the same size, which is equal to $\lfloor C/M \rfloor$.

Implementation Details Our experiments use evaluation scripts derived from the EleutherAI Language Model Evaluation Harness (Gao et al., 2023).² To quantize the models we use scripts from Auto-GPTQ package.³ We run quantization and inference for all the experiments on a single NVIDIA A-100 GPU. For the largest model, uncompressed OPT-13B, the evaluation run took roughly two hours for all the datasets. Frantar et al., 2022 report GPTQ runtime for the models.

²<https://github.com/EleutherAI/lm-evaluation-harness>

³<https://github.com/PanQiWei/AutoGPTQ>

C Confidence Evaluation in LLMs after Quantization

In this last experiment we study the evolution of the confidence score for our different models on the six studied datasets. More precisely, we study the mean difference of confidence score between full and quantized models for different ranges of confidence scores of the full model.

As presented in Figure 3, the change of probabilities is the lowest one when the model is over-confident and the uncertainty of the model is impacted (*i.e.*, increased) by the quantization. This observation goes hand in hand with the entropy values, serving as a measure of model uncertainty, shown in Table 5 (Appendix D). We also note that, in the case of binary problems (PIQA, BoolQ and XSTORY CLOSE EN), that the most impacted confidence scores are the ones for which the model is not confident in its prediction.




Figure 3: Confidence difference for models across datasets. For each dataset (in column) and each model (in line), we provide the difference in prediction scores between the full and quantized models. More precisely, each bar represents the mean difference in confidence between the quantized and full models, with confidence in the full model represented on the horizontal axis. Note that some ranges start from 0.5 for binary tasks and 0.25 for multi-class (with four classes) tasks. For a confidence lower than the previous one, there is no chance of being assigned to the associated class.

D Confidence Evaluation Results









Model	ARC EASY		BOOLQ		HELLASWAG		OPENBOOKQA		PIQA		XSTORY	
	Conf.	Conf _{err}	Conf.	Conf _{err}	Conf.	Conf _{err}	Conf.	Conf _{err}	Conf.	Conf _{err}	Conf.	Conf _{err}
 ...												
7B	88.45	75.44	76.21	64.8	96.85	95.02	79.2	78.93	94.78	89.39	95.23	95.57
7BQ	88.15	75.61	76.75*	65.56	96.89	95.13	79.01	79.07	94.67	88.8	95.26	95.72
 ...												
7B	84.68	73.12	75.75	67.78	96.8	95.34	78.64	78.85	94.24	90.32	95.01	95.51
7BQ	81.6*	71.93	68.95*	63.55	96.48*	95.13	78.25	78.68	93.99	90.05	94.8	94.94
 ...												
560M	76.45	73.76	64.74	64.38	96.39	96.27	78.28	78.77	91.47	89.8	94.47	94.62
560MQ	75.89	73.68	61.89*	62.76	96.47	96.35	78.74	79.11	91.76	90.07	94.55	95.22
1.1B	76.2	72.22	70.63	69.16	96.45	96.16	78.3	78.81	92.1	89.99	94.36	94.46
1.1BQ	76.0	72.95	73.28*	72.08	96.52	96.18	77.97	78.31	91.78	89.58	94.21	94.87
1.7B	77.47	72.86	76.12	74.9	96.24	95.89	78.05	78.57	91.89	89.75	93.96	94.08
1.7BQ	76.34*	72.44	76.06	74.85	96.11	95.91	77.52	78.31	91.54	89.56	93.87	94.32
3B	78.6	73.11	72.5	70.55	96.24	95.75	78.3	78.56	92.37	89.57	94.36	94.08
3BQ	77.24*	72.23	71.59*	69.93	96.43	96.02	77.5	77.9	92.25	89.2	94.15	94.32
7.1B	79.95	73.11	69.97	66.55	96.26	95.64	78.36	78.52	92.71	88.05	94.59	94.85
7.1BQ	79.46	72.85	69.59*	66.58	96.3	95.62	78.17	78.59	92.53	89.03	94.56	94.54
 ...												
125M	75.9	74.29	67.95*	67.42	96.31	96.29	77.6	78.42	90.96	89.31	94.1	94.86
125MQ	75.88	74.57	64.45	64.21	96.29	96.15	78.14	79.62	91.39	89.89	94.31	94.59
350M	75.45	73.25	67.45	66.57	96.07	95.91	78.36	79.38	91.39	88.85	94.17	94.62
350MQ	76.46*	74.84	63.03*	62.34	96.25	96.03	78.26	78.68	91.33	89.23	94.38	94.43
1.3B	77.67	72.44	64.25	62.24	96.31	95.74	78.35	79.07	91.91	88.69	94.39	94.47
1.3BQ	77.02*	72.54	64.26	63.5	96.14	95.58	78.56	79.16	91.94	88.59	94.31	94.87
2.7B	78.22	71.73	63.67	61.81	96.32	95.51	78.45	78.68	91.89	87.93	94.42	94.6
2.7BQ	77.58*	71.69	63.66	62.07	96.2	95.62	77.89	77.66	92.14	88.02	94.35	94.64
6.7B	80.46	72.14	65.88	62.46	96.51	95.57	78.65	79.16	93.29	89.78	94.38	95.32
6.7BQ	80.29	72.52	64.16*	60.9	96.5	95.55	78.66	78.32	93.13	89.4	94.55	94.69
13B	81.36	72.42	67.3	63.32	96.49	95.48	78.7	78.75	93.23	88.64	94.98	95.53
13BQ	80.96*	72.4	66.78*	62.35	96.5	95.52	79.08	79.46	93.03	88.77	94.77	95.18

Table 4: Mean confidence evaluation results across benchmarks. Conf.: Mean confidence in predictions; Conf_{err}: Mean confidence in wrong predictions. The * is used to denote a significant difference with a confidence level set at 0.05 (paired *t*-test). Q denotes quantized models. Notations: =MISTRAL; =LLAMA; =BLOOM; =OPT.









Model	ARC EASY		BOOLQ		HELLASWAG		OPENBOOKQA		PIQA		XSTORY	
	Conf _{true}	H	Conf _{true}	H	Conf _{true}	H	Conf _{true}	H	Conf _{true}	H	Conf _{true}	H
 ...												
7B	76.82	43.09	71.37	71.05	61.14	10.96	30.93	72.41	79.52	17.69	47.4	16.05
7BQ	75.76*	44.42*	71.37	70.07*	59.73*	10.87	30.36	73.19	79.14	18.11	47.27	16.1
 ...												
7B	70.22	56.26	66.83	71.19	56.83	11.37	31.28	73.87	77.04	19.53	46.88	16.62
7BQ	65.01*	66.4*	61.48*	84.3*	53.69*	12.21*	28.42*	76.09	75.3*	20.61*	46.97	17.65*
 ...												
560M	44.01	82.38	51.83	90.77	31.5	12.41	17.68	75.57	62.89	27.68	47.93	18.72
560MQ	42.42*	83.81*	49.78*	93.37*	31.07*	12.26	17.65	73.97	62.07	27.52	48.16	18.61
1.1B	47.33	83.0	54.95	82.92	34.51	12.35	19.62	75.19	65.82	25.78	48.11	19.33
1.1BQ	45.23*	83.28	55.5*	79.2*	33.67*	12.28	18.26*	76.23	64.99*	26.44	48.3	19.66
1.7B	51.15	79.99	57.09	74.65	37.52	13.11	21.19	75.97	67.06	26.28	46.93	19.99
1.7BQ	49.21*	83.58*	56.95	74.95	36.52*	13.38	19.83*	77.84*	66.23*	27.16*	47.41	20.4
3B	54.29	76.14	56.73	79.88	41.26	12.91	22.3	74.9	69.3	25.11	47.1	19.1
3BQ	52.07*	80.42*	56.34*	81.98*	40.41*	12.55*	21.64	77.03*	68.75	25.6	47.57	19.55
7.1B	59.1	71.97	57.67	83.21	46.24	12.87	24.64	75.42	71.88	23.84	46.8	18.6
7.1BQ	57.79*	73.52*	57.33*	83.89*	45.23*	12.89	23.74*	75.9	71.38	24.41*	46.69	18.65
 ...												
125M	40.09	83.61	52.46	86.7	29.04	12.86	16.95	77.19	61.88	28.97*	48.26	20.12
125MQ	39.12*	83.58	50.99*	91.0*	28.62*	12.75	16.27	76.26	61.22	28.24	48.33	19.54
350M	41.01	85.36	53.4	87.14	32.03	13.51	17.41	76.08	63.78	28.0	47.69	19.77
350MQ	40.69	82.59*	51.72*	92.36*	31.84	13.09	17.05	75.09	62.73*	28.31	48.04	19.16
1.3B	52.33	79.37	53.88	90.87	41.5	13.03	22.51	75.63	70.02	26.08	47.12	18.73
1.3BQ	50.83*	81.23*	52.07*	91.22	40.47*	13.26	22.2	75.09	69.52	25.96	47.21	18.76
2.7B	55.63	77.64	54.34	91.63	45.84	12.84	24.63	74.97	72.0	25.68	46.89	18.82
2.7BQ	53.91*	79.47*	52.88*	91.66	45.12*	13.09	23.71	76.5	71.45	25.45	46.9	18.78
6.7B	60.58	70.14	57.42	88.65	50.37	12.12	26.42	74.26	74.42	22.67	46.81	18.3
6.7BQ	60.08	70.89	56.57*	90.78*	49.78*	12.22	26.45	74.65	74.26	22.96	46.63	18.06
13B	62.35	67.09	58.24	86.84	52.15	12.25	26.85	73.83	74.56	22.26	46.88	17.13
13BQ	62.06	68.26*	58.38	87.45*	51.61*	12.24	26.67	73.7	74.38	22.69*	46.98	17.38

Table 5: Mean confidence in true classes and predictive entropy evaluation results across benchmarks. Conf_{true}: Mean confidence in true class; H=Mean Predictive entropy in the answers, multiplied by 100. The * is used to denote a significant difference with a confidence level set at 0.05 (paired *t*-test). Q denotes quantized models. Notations: =MISTRAL; =LLAMA; =BLOOM; =OPT.