

Cloe: Hybrid Surgical Distillation for 1.58-Bit Edge Quantization of Large Language Models

Anirudh Malik, Vishnuprasad, Poojith Devan, M Sparsh Mehra
OneBit AI

v1.0 Production Ready

Abstract

The deployment of Large Language Models (LLMs) on edge hardware is severely bottlenecked by memory footprint and computational requirements. Extreme quantization approaches, such as 1.58-bit (ternary) weights, offer a promising solution by mapping network parameters to $\{-1, 0, 1\}$. However, naive Post-Training Quantization (PTQ) applied to sub-1B parameter models consistently results in catastrophic degradation due to a lack of parameter redundancy. In this paper, we present **Cloe**, an empirical investigation into extreme quantization targeting edge GPUs (e.g., NVIDIA RTX 3050 4GB). After detailing five failed methodologies including naive quantization and layerwise activation matching, we introduce **Hybrid Surgical Distillation**. By observing that attention layers correspond to logical reasoning while Multilayer Perceptrons (MLPs) act as factual memory banks, we surgically quantize only the MLP layers to 1.58-bits while preserving Attention layers in 16-bit (BFLOAT16). Utilizing knowledge distillation to recover representational integrity, Cloe v1 achieves 44.00% on MMLU, outperforming standard 16-bit baseline models in its class while exhibiting a 44% higher intelligence density per gigabyte of VRAM.

1 Introduction

Large Language Models (LLMs) have demonstrated unprecedented capabilities, yet their massive parameter counts restrict their usage to expensive, cloud-based GPU clusters. A “small” 0.8B parameter model necessitates approximately 1.6 GB of VRAM in standard 16-bit precision, while “massive” 70B models require over 140 GB. This memory wall precludes the deployment of localized AI on consumer’s edge devices.

The primary objective of Cloe is to aggressively compress an LLM such that its weights map to a ternary state $\{-1, 0, +1\}$. This 1.58-bit quantization ($\log_2(3) \approx 1.58$) yields a roughly 10× compression ratio over FP16 matrices and replaces computationally expensive floating point multiplications with efficient additions and subtractions.

However, while extreme quantization is feasible on highly redundant 70B models, applying it to compact architectures like Qwen 3.5 0.8B [1] leads to structural collapse. We detail our iterative journey through layer-wise repairs and full-network distillation, ultimately culminating in a hybrid architecture that successfully bridges the gap between efficiency and cognitive capability.

2 Background & Formulations

2.1 Base Architecture

Our base model is Qwen 3.5 0.8B, a bilingual (Chinese-English) autoregressive transformer consisting of 24 layers. Each block contains two primary sub-structures: Multi-Head Attention (MHA) and a Feed-Forward Network (MLP). The standard formulation relies on BFloat16 ('torch.bfloat16') for weight representation.

2.2 The 1.58-Bit Quantization Paradigm

Unlike standard integer quantization (e.g., INT8 or INT4), ternary quantization forces weights into three absolute states. For a given weight matrix W , we compute a floating-point scaling factor γ based on the mean absolute value of the layer:

$$\gamma = \frac{1}{n} \sum_{i=1}^n |W_i| \quad (1)$$

The quantized matrix $\widetilde{W}$ is then obtained via:

$$\widetilde{W} = \text{Round} \left(\text{Clamp} \left(\frac{W}{\gamma + \epsilon}, -1, 1 \right) \right) \times \gamma \quad (2)$$

where ϵ is a small constant (e.g., $1e^{-9}$) for numerical stability.

3 The Core Invention: BitLinear

To enable gradient flow through the non-differentiable $\text{Round}(\cdot)$ operation during training, we implement a **BitLinear** module utilizing a Straight-Through Estimator (STE).

Algorithm 1 BitLinear Forward Pass with STE

Require: Input X , Full-precision Weights W , Bias B

- 1: $\gamma \leftarrow \text{mean}(|W|)$
 - 2: $W_{\text{scaled}} \leftarrow \text{clamp}(W/(\gamma + \epsilon), -1, 1)$
 - 3: $W_{\text{quant}} \leftarrow \text{round}(W_{\text{scaled}}) \times \gamma$ ▷ Straight-Through Estimator Trick
 - 4: $\widetilde{W} \leftarrow W + (W_{\text{quant}} - W).\text{detach}()$
 - 5: $Y \leftarrow X\widetilde{W}^T + B$
 - 6: **return** Y
-

In earlier iterations, we incorporated an additional RMSNorm pre-activation step to prevent magnitude collapse and a λ -ramp schedule to ease the transition from FP16 to 1.58-bit. However, our final hybrid methodology allowed for the simplified module shown in Algorithm 1.

4 The Path to Cloe: Experimental Phases

Developing an edge-capable 1.58-bit model required navigating several catastrophic failures induced by compounding quantization errors.

4.1 Phase 1: Naive Full Quantization (Failed)

Replacing all ‘nn.Linear’ layers with ‘BitLinear’ and training via standard Causal Language Modeling resulted in complete linguistic degradation. The model output non-linguistic noise (e.g., “ $\Delta\Delta\Delta rksjf223\dots$ ”). **Insight:** Sub-1B models lack the parameter redundancy required to survive global ternary mapping.

4.2 Phase 2 & 3: Layer-Wise & Distillation (Partial)

We attempted to repair the model by freezing a BFloat16 “Teacher” and training a quantized “Student” to minimize Mean Squared Error (MSE) on intermediate layer activations.

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum ||h_S^{(i)} - h_T^{(i)}||_2^2 \quad (3)$$

Despite individual layers matching closely, micro-errors compounded exponentially over 24 layers. Transitioning to output Knowledge Distillation (Phase 3) decreased the loss but failed to recover grammatical coherence.

4.3 Phase 4 & 5: Master Reset & Voice Recovery (Partial)

We initialized the model with perfect Teacher weights and selectively quantized while keeping the Language Modeling (LM) Head unfrozen. Later, we froze the entire quantized body and trained *only* the LM head to map damaged hidden states to a valid vocabulary space. While the model began producing real words, logical reasoning was entirely absent.

5 The Breakthrough: Hybrid Surgical Distillation

The critical pivot was identifying the distinct roles of transformer sub-layers. **Attention layers govern contextual logic and long-range dependencies, while MLPs act as factual memory repositories.** Memory is highly compressible; logic is brittle.

5.1 The Hybrid Architecture

We established a strict dichotomy within the transformer blocks:

- **Attention Layers:** Preserved in full 16-bit precision (BFloat16) to maintain reasoning capabilities.
- **MLP Layers:** Surgically replaced with 1.58-bit BitLinear modules.




Figure 1: Hybrid Surgical Architecture of Cloe v1.

5.2 Surgical Distillation Optimization

With the architectural damage contained entirely within the MLP pathway, we employed Kullback-Leibler (KL) Divergence distillation over 400 high-quality samples from the **FineWeb-Edu** dataset. The student network was trained to perfectly mimic the probability distribution of the 16-bit teacher.

The loss function utilized temperature scaling ($\tau = 2.0$) to soften the logits:

$$\mathcal{L}_{KD} = \tau^2 \sum_{k=1}^V P_T(k|\tau) \log \left(\frac{P_T(k|\tau)}{P_S(k|\tau)} \right) \quad (4)$$

where P_T and P_S are the softmax probabilities of the Teacher and Student respectively, and V is the vocabulary size. Upon convergence, weights were “baked” permanently into matrices of $\{-1, 0, 1\}$.

6 Evaluation & Benchmarks

We evaluated `cloe_v1_master.pt` utilizing an `EvalScope`-based suite and custom fast-evaluation scripts targeting edge GPUs.

Table 1: Comparative Benchmarks: Cloe v1 vs. Baseline Qwen Models

Model	Precision	MMLU (%)	GSM8K (%)	Disk	VRAM	Intell. Density
Standard Qwen 0.5B	16-bit	39.50	~0.0	1.1 GB	~1.5 GB	26.3 pts/GB
Standard Qwen 1B	16-bit	42.00	8.0	2.2 GB	~3.5 GB	19.1 pts/GB
Cloe v1 (Ours)	Hybrid 1.58-bit	44.00	5.0	1.6 GB	~2.0 GB	27.5 pts/GB

6.1 Metrics and Results

As shown in Table 1, Cloe v1 achieves a **44.00%** on MMLU (Measuring Massive Multitask Language Understanding) and **5.00%** on the mathematically rigorous GSM8K dataset.

Astonishingly, despite housing ternary weights in the MLP modules, Cloe v1 outperforms both the 0.5B and 1B standard Qwen baselines in MMLU score. We introduce the metric *Intelligence Density* (MMLU Score / Peak VRAM GB). Cloe yields an intelligence density of 27.5 pts/GB, which is a **44% improvement** in spatial efficiency compared to the standard 1B model (19.1 pts/GB).

6.2 Inference Acceleration

For production inference, no custom `BitLinear` classes are required in memory. Because the baked ternary weights are natively stored in the model’s `state_dict`, standard `nn.Linear` matrices dynamically multiply the $\{-1, 0, 1\}$ values by their respective pre-calculated γ scales. This execution yields significantly lower latency and cache-miss rates on the target RTX 3050 hardware.

7 Conclusion

Project Cloe demonstrates that the dream of efficient, edge-native LLMs is achievable without waiting for entirely new pre-training runs from massive corporations. By mapping transformer Attention to “logic” (preservation) and MLP to “memory” (compression), our Hybrid Surgical Distillation methodology successfully stabilized a 1.58-bit sub-1B architecture. Cloe v1 establishes a new paradigm for extracting maximum linguistic and reasoning density out of constrained edge hardware.

References

- [1] Bai, Jinze, et al. “Qwen technical report.” *arXiv preprint arXiv:2309.16609* (2023).
- [2] Ma, Shuming, et al. “The era of 1-bit llms: All large language models are in 1.58 bits.” *arXiv preprint arXiv:2402.17764* (2024).
- [3] Hinton, Geoffrey, Oriol Vinyals, and Jeff Dean. “Distilling the knowledge in a neural network.” *arXiv preprint arXiv:1503.02531* (2015).

A Glossary of Key Terms

- **1.58-bit:** Derived from $\log_2(3)$, representing the information content needed to store exactly three states ($\{-1, 0, +1\}$).
- **STE (Straight-Through Estimator):** A mathematical proxy utilized during backpropagation to assign a gradient of 1 to non-differentiable step/rounding functions.
- **Baking:** The final post-training process where FP16 proxy weights are permanently overwritten with their deterministic ternary output and scale factors.
- **Intelligence Density:** A proposed efficiency metric representing Benchmark Score divided by maximum runtime VRAM footprint.