

A Compute Resources and Cost

We utilize our internal cluster equipped with A-100 GPUs, each with 80GB of memory. Each unsupervised fine-tuning phase for a 7B model was conducted on a single node. As detailed in Section 3.2, we perform up to 20 epochs on a 75M-token dataset, which takes approximately half a day. This process is repeated for three different types of augmentations in the training data, as described in Section 4, as well as for one Llama1-7B, one Llama2-13B, and one Llama3-8B post-training. In total, this amounts to approximately $12 \times 8 \times 6 \approx 600$ GPU hours. All subsequent SFT and unlearning phases are more than 10 times less resource-intensive (see Section 3.2), so we omit them here. Considering all failed experiments and those not included in the paper, we estimate that the total training time was close to 2000 GPU hours.