


GPU benchmark (device=cpu)

(a) Training throughput



(b) Inference latency

