Heavy-tailed update distributions arise from information-driven self-organization in nonequilibrium learning

Anonymous Submission for Network Science Informs AI (NSIA) 2026

Like human decision-making under real-world constraints, artificial neural networks may balance free exploration in parameter space with task-relevant adaptation. In this study, we identify consistent signatures of criticality during neural network training and provide theoretical evidence that such scaling behaviour arises naturally from information-driven self-organization: a dynamic balance between the maximum entropy principle that promotes unbiased exploration, and mutual information constraint that relates updates with task objective. We numerically demonstrate that the power-law exponent of updates remains stable throughout training, supporting the presence of self-organized criticality. Furthermore, we show that the loss landscape exhibits exponential ruggedness under small perturbations, transitioning to power-law ruggedness at larger scales, in the absence of mini-batch noise, indicating an intrinsic geometric landscape. We also observe a power-law distribution in the intervals between large updates, indicating an intermittent learning process. Together, these findings suggest that neural network learning reflects a nonequilibrium process governed by the fundamental trade-off between randomness and relevance, highlighting its dynamic nature and offering insights into the interpretability of artificial intelligence systems.

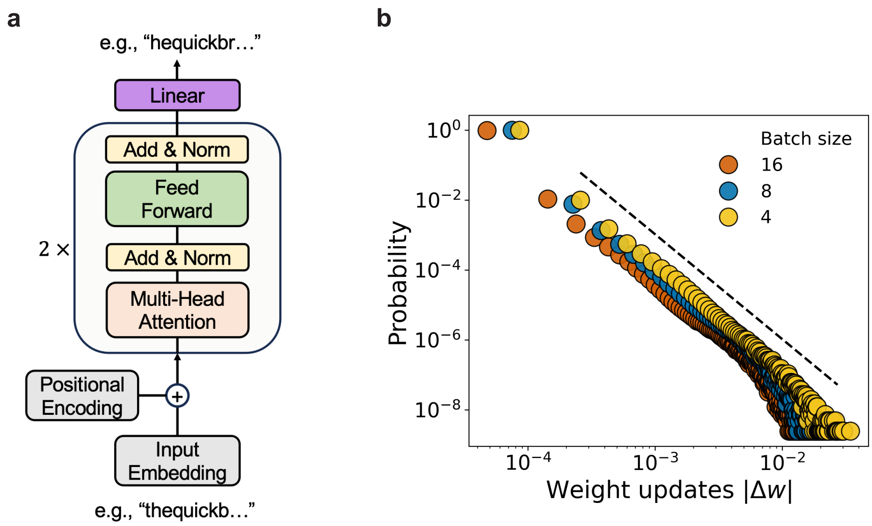

Figure 1: Character-level Transformer model and heavy-tailed updates. (a) Schematic of the simplified Transformer encoder used in our analysis. (b) Empirical distribution of weight update magnitudes $|\Delta w|$ during training using stochastic gradient descent (SGD) with learning rate $r=0.1$.