The Shadow Price of Reasoning:
Economic Perspective on Optimal Budget Allocation for LLMs

Xu Wan · Speed Zhu · Jianwei Cai · Guang Chen · XiMing Huang · Wiggin Zhou · Mingyang Sun (Corr. Author)

ICML International Conference On Machine Learning

Tencent Hunyuan

Motivation

Reasoning tasks differ not only in difficulty, but also in where additional tokens become useful.

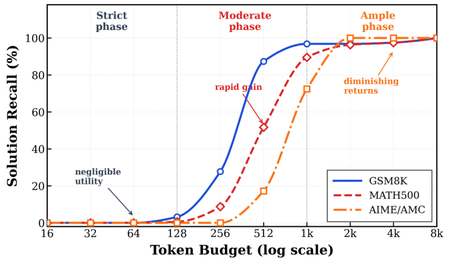

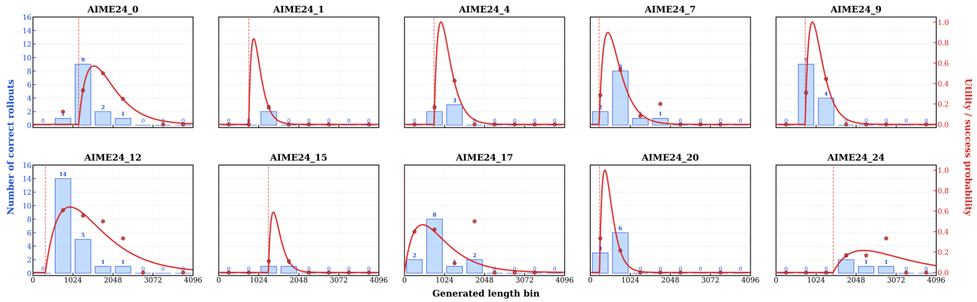

Method

Treat tokens as a global budget and allocate them where marginal utility is highest.

◆ **Model** token utility with a shifted-surge curve.

$$\phi_i(t) = \begin{cases} 0 & 0 \le t < \tau_i \\ \alpha_i (t - \tau_i) \cdot e^{-\beta_i (t - \tau_i)} & t \ge \tau_i \end{cases} \quad \Rightarrow \quad t_i^*(\lambda) = \tau_i + \frac{1}{\beta_i} \left[1 - W_0 \left(\frac{\lambda e}{\alpha_i} \right) \right]$$

◆ **Predict** each query's emergence threshold.

$$\hat{\tau}_i = \exp(f_\theta(s_i)).$$

Individual Optimal Allocation Policy

◆ **Solve** a global shadow price to allocate tokens, abandoning queries with negative surplus.

$$\max_{\mathbf{t} \in \mathbb{R}^N_{\ge 0}} \sum_{i=1}^{N} \phi_i(t_i) \quad \text{s.t.} \quad \sum_{i=1}^{N} t_i \le B_{\text{total}}.$$

Results

CLEAR improves over uniform by up to +24.0 points by reallocating tokens from low-surplus queries to solvable ones.

Method	Balanced Stream				Mostly-Easy Stream				Mostly-Hard Stream				U-Shaped Stream			
	256	512	1024	2048	256	512	1024	2048	256	512	1024	2048	256	512	1024	2048
Uniform	3.0	12.4	17.4	21.6	9.0	33.8	45.6	**49.0**	1.0	5.0	6.8	11.4	4.4	18.4	25.4	27.2
Predictor	0.6	4.6	**18.8**	**22.8**	1.8	17.6	46.0	**49.0**	0.0	1.6	8.2	**12.8**	0.8	9.0	26.4	27.6
TALE-EP	3.8	6.8	16.0	22.2	7.2	10.8	22.4	39.8	4.2	5.8	**9.2**	12.0	7.2	10.4	24.6	**29.0**
CLEAR (Heaur.)	11.0	**17.0**	**18.8**	**22.8**	31.6	37.8	46.0	**49.0**	4.0	**7.2**	8.2	**12.8**	17.8	20.6	26.4	27.6
CLEAR (Auction)	14.2	15.6	**18.8**	**22.8**	32.8	39.2	46.0	**49.0**	6.0	6.6	8.2	**12.8**	18.4	22.2	26.4	27.6
CLEAR (Lambert)	**14.6**	16.4	**18.8**	**22.8**	**33.0**	**40.4**	**47.6**	**49.0**	**6.2**	6.8	8.2	**12.8**	**18.6**	**23.2**	**26.8**	27.6
Oracle	19.0	21.6	23.6	26.0	42.2	48.4	49.8	49.8	9.6	11.4	12.8	16.6	26.4	27.6	30.2	32.8
Gain (Abs.)	**+11.6**	+4.0	+1.4	+1.2	**+24.0**	+6.6	+2.0	+0.0	**+5.2**	+1.8	+1.4	+1.4	**+14.2**	+4.8	+1.4	+0.4