


The Illusion of Stochasticity in LLMs

Xiangming Gu, Soham De, Michalis Titsias, Larisa Markeeva, Petar Veličković, Razvan Pascanu

Necessity of reliable independent sampling




LLMs are prompted to generate MCQs with bias to “C” as answers

Stackelberg security game:
LLM agents are prompted to choose one server to inspect

Model	A	B	C	D	Best-response attacker success
Gemini-3 Flash	0.034	0.521	0.373	0.071	0.966
Gemini-3.1 Pro	0.002	0.297	0.700	0.001	0.999
Gemini-3.5 Flash	0.001	0.408	0.571	0.020	0.999

LLMs’ illusion of stochasticity can be **exploited**

Failures of LLMs in independent sampling from distributions







Logit analysis on pre-trained/post-trained LLMs shows **post-trained** LLMs have stronger biases

Failures happen regardless of LLM families, distributions, decoding parameters, COT

Sequential sampling

With last/all history samples as context, then sample next one




Batched sampling

Prompt to sample multiple values

Sequential sampling introduces **repulsive** temporal bias
Batched sampling introduces **periodic** temporal bias

auto-correlation


Tool-use PRNG



Without extra prompting

Prompting to use current time as **seed**

Distribution conversion



U[0, 1] seeds the LLM inputs