

# The Effect of Sampling Temperature on Problem Solving in Large Language Models

**Matthew Renze**  
Johns Hopkins University  
mrenze1@jhu.edu

**Erhan Guven**  
Johns Hopkins University  
eguwen2@jhu.edu

## Abstract

In this research study, we empirically investigate the effect of sampling temperature on the performance of Large Language Models (LLMs) on various problem-solving tasks. We created a multiple-choice question-and-answer (MCQA) exam by randomly sampling problems from standard LLM benchmarks. Then, we used nine popular LLMs with five prompt-engineering techniques to solve the MCQA problems while increasing the sampling temperature from 0.0 to 1.6. Despite anecdotal reports to the contrary, our empirical results indicate that changes in temperature from 0.0 to 1.0 do not have a statistically significant impact on LLM performance for problem-solving tasks. In addition, these results appear to generalize across LLMs, prompt-engineering techniques, and problem domains. All code, data, and supplemental materials are available on GitHub at: <https://github.com/matthewrenze/jhu-llm-temperature>.

## 1 Introduction

### 1.1 Background

In recent years, Large Language Models (LLMs) have revolutionized the field of artificial intelligence. The availability of open-source LLMs and pay-per-use APIs has allowed engineers to incorporate LLMs in their AI systems. However, prompt engineering and hyperparameter tuning are required to work effectively with LLMs.

Prompt-engineering techniques help LLMs solve complex problems, avoid hallucinations, and provide more accurate responses. For example, we can use techniques like chain-of-thought, tree-of-thought, self-criticism, and self-consistency to improve LLM performance (Mialon et al., 2023; White et al., 2023).

In addition, several inference hyperparameters can be adjusted to modify the LLM's output at runtime. For example, hyperparameters like sampling

temperature, top-k sampling, repetition penalty, and maximum token length all affect the LLM's output and performance (OpenAI, 2023d; Llama-2-Team, 2023; Wang et al., 2023a).

Despite significant interest in LLMs and progress in LLM best practices, many open questions remain about optimal prompt-engineering techniques and inference hyperparameters for LLMs. To complicate matters, various local optima may exist for LLMs, prompt types, and problem domains (Wang et al., 2023a).

The prompt-engineering community has an abundance of opinions and anecdotal evidence regarding optimal prompt-engineering techniques and inference hyperparameter settings. However, we currently lack systematic studies and empirical evidence to support many of these claims.

As a result, this paper aims to address the open question of the optimal LLM sampling temperature for problem-solving tasks. In addition, we aim to provide a systematic study with empirical results to add to the growing body of knowledge used to create LLM and prompt-engineering best practices.

### 1.2 Sampling Temperature

Sampling temperature is a hyperparameter of an LLM used in a temperature-based sampling process. It controls the randomness of the model's output at inference time (Ackley et al., 1985; Hinton et al., 2015; Wang et al., 2020, 2023a).

During each step of an LLM's decoding process, the LLM uses the previous tokens to choose the next output token. The final layer of the LLM uses a softmax function to convert raw scores (logits) into probabilities.

In greedy sampling, the model will always choose the most likely next token. However, for probabilistic sampling, the next token is selected from a probability distribution.

Temperature sampling is a modification to the softmax function, which adjusts the resulting prob-

ability mass functions. In this modified softmax function,  $v_k$  is the  $k$ -th vocabulary token,  $l_k$  is the token’s logit, and  $\tau$  is a constant temperature. See equation 1.

$$\Pr(v_k) = \frac{e^{l_k/\tau}}{\sum_i e^{l_i/\tau}} \quad (1)$$

A lower temperature makes the output of the LLM more deterministic, thus favoring the most likely predictions. This conservativeness is captured by the model’s tendency to produce more repetitive, focused, and less diverse output based on the patterns most commonly seen in the training data (Hinton et al., 2015; Wang et al., 2020, 2023a).

A higher temperature increases the randomness of the output, thus favoring more “creative” predictions. This creativity is captured by the model’s willingness to explore more unconventional and less likely outputs. Higher temperatures can lead to novel text, diverse ideas, and creative solutions to problems (Hinton et al., 2015; Wang et al., 2020, 2023a).

In the context of problem-solving, temperature can be seen as a trade-off between exploring and exploiting possible solutions within the solution space. Lower temperatures tend to exploit more probable solutions; higher temperatures explore the solution space more broadly.

### 1.3 Choosing a Sampling Temperature

Within the prompt-engineering community, there are a variety of opinions and best practices regarding the ideal sampling temperature for various problem-solving tasks (Microsoft, 2023; Shieh, 2024).

Low sampling temperatures are recommended for tasks requiring precision and factual accuracy, such as technical writing, code generation, or question-answering (Xu et al., 2022; Zhu et al., 2023). However, higher temperatures are recommended for tasks requiring creativity, such as writing poetry, creating stories, or brainstorming.

Higher temperatures also increase the probability of model hallucination. Hallucination is a phenomenon where an LLM produces statistically probable responses that are factually incorrect or nonsensical. As a result, optimal temperature selection is also a balance between creativity and hallucination (Lee, 2023).

Practical guidelines for choosing a sampling temperature for a specific task or problem domain are often vague or anecdotal. Prompt-engineering

guides often provide hypothetical examples of optimal sampling temperatures for various tasks. However, they rarely cite any sources or provide empirical evidence.<sup>1</sup>

As a result, the current state of choosing the optimal sampling temperature for specific problems is largely based on guesswork, gut instinct, non-systematic experimentation, and iterative refinement.<sup>2,3</sup>

## 2 Methods

### 2.1 Models

The models used in this research project comprise nine widely-used foundational LLMs. To complement our analysis, we also conducted experiments using five prompts created using commonly used prompt-engineering techniques.

First, we reviewed the prior literature to identify candidate LLMs commonly used for problem-solving tasks. We limited our candidate models to those that allowed the model’s sampling temperature to be specified via their API (OpenAI, 2022a, 2023b,c; Llama-2-Team, 2023). See Table 1 for a list of LLMs used in the experiment.

Next, we reviewed the existing literature for commonly used prompt-engineering techniques. We limited our candidate prompts to those that could be performed in a single request-and-response cycle with one-shot in-context learning. We excluded multi-step agents, few-shot learning, and model fine-tuning.

As a result, we selected five prompt-engineering techniques to construct our system prompts:

- **Baseline** - no prompt engineering; the LLM is instructed to return only a single multiple-choice answer as its output (e.g., ‘Answer(“C”)’).
- **Domain Expertise** – the system prompt specifies that the LLM is an expert in the problem

<sup>1</sup>A few empirical studies exist that indicate sampling temperature does have an effect on LLM performance on some types of problem-solving tasks (e.g., code generation, engineering exams, etc.) (Xu et al., 2022; Pursnani et al., 2023; Zhu et al., 2023).

<sup>2</sup>For example, OpenAI’s GPT-3.5 API allowed users to set the sampling temperature from 0.0 to 1.0 with a default of 0.7. GPT-4’s API expanded this range from 0.0 to 2.0 with a default of 1.0. No explanation from OpenAI has been provided for these default values or their change from GPT-3.5 to GPT-4 (OpenAI, 2023a).

<sup>3</sup>Even the GPT-4 Technical Report explains that the authors used their “best-guess” when choosing sampling temperatures while evaluating GPT-4 on various benchmarks. See Appendix A in the GPT-4 Technical Report (OpenAI, 2023c).

Name	Vendor	Released	License	Source
Claude 3 Opus	Anthropic	2024-03-04	Closed	(Anthropic, 2024b,a)
Command R+	Cohere	2024-04-04	Open	(Cohere, 2024a,b)
Gemini 1.0 Pro	Google	2023-12-06	Closed	(Pichai and Hassabis, 2023; Gemini-Team, 2023)
Gemini 1.5 Pro (Preview)	Google	2024-02-15	Closed	(Pichai and Hassabis, 2024; Gemini-Team, 2024)
GPT-3.5 Turbo	OpenAI	2022-11-30	Closed	(OpenAI, 2022a,b)
GPT-4	OpenAI	2023-03-14	Closed	(OpenAI, 2023b,c)
Llama 2 7B Chat	Meta	2023-07-18	Open	(Meta, 2023; Llama-2-Team, 2023)
Llama 2 70B Chat	Meta	2023-07-18	Open	(Meta, 2023; Llama-2-Team, 2023)
Mistral Large	Mistral AI	2024-02-26	Closed	(Mistral-AI-Team, 2024)

Table 1: LLMs used in the experiment.

domain of the exam (e.g., “medicine”) or the topic of the problem (e.g., “anatomy”) (White et al., 2023).

- **Self-recitation** – the system prompt instructs the LLM to recite its own internal knowledge about the problem before answering the question (Sun et al., 2023; White et al., 2023).
- **Chain-of-Thought (CoT)** – the system prompt instructs the LLM to “think step-by-step” to encourage it to reason through the problem procedurally (Kojima et al., 2022; Wei et al., 2022).
- **Composite** – the system prompt combines domain expertise, self-recitation, chain-of-thought, and adds self-criticism (Huo et al., 2023; Wang et al., 2023b).

Finally, we provided the LLM with a single example problem-and-solution pair for one-shot in-context learning. The example solution was adapted for each prompt based on the prompt-engineering technique used. For example, the CoT prompt included a chain of thought in its solution. See Figure 10 in the Appendix for a sample prompt.

## 2.2 Data

The test dataset used in this research study consists of a series of Multiple-Choice Question-and-Answer (MCQA) exams derived from widely used LLM performance benchmarks.

First, we reviewed the prior literature to identify benchmarks frequently used to evaluate LLMs. We limited our candidate benchmarks to those containing MCQA problems so that we could use correct-answer accuracy as our primary performance metric.

Next, we selected a set of problems that covered a range of problem domains (e.g., math, science, law, etc.) and difficulty levels (e.g., secondary

school, university, etc.) These problem sets can be seen in Table 2.

Then, we converted the benchmark problems from their original data format into a standardized data structure using the JSON Lines (JSON-L) format (Ward, 2014). Our standardized set of exams allowed us to use the exams interchangeably without modifying the code in the test harness. See Figure 11 in the Appendix for a sample of an MCQA problem.

Finally, we created two MCQA exams of different sizes. We created a large exam with 1,000 questions by randomly sampling 100 problems from each of the ten problem sets. This 1,000-question (large) exam was used with GPT-3.5 to perform a detailed analysis of temperature across problem domains.

Additionally, we created a smaller exam of 100 questions by randomly sampling ten questions from each of the ten domain-specific problem sets. This 100-question (small) exam was used for our high-level analysis of sampling temperature across all nine models, all five prompt-engineering techniques, and extended temperature range (0.0-1.6).<sup>4</sup>

## 2.3 Process

Our experiment was designed to test the problem-solving performance of LLMs across ten models, five prompt-engineering techniques, ten problem domains, 100 problems within each problem domain, and all viable sampling temperatures. For each combination of model, prompt, exam, and temperature, we instructed the LLM to answer each question ten times so we could assess the average correct-answer accuracy.

The full experiment setup can be seen in Figure 1 and Algorithm 1. However, due to cost and runtime considerations, we conducted a subset of the full experiment designed to capture the most valuable

<sup>4</sup>We used the smaller 100-question exam due to cost and runtime considerations.

Problem Set	Benchmark	Domain	Questions	License	Source
ARC Challenge Test	ARC	Science	1,173	CC BY-SA	(Clark et al., 2018)
AQUA-RAT	AGI Eval	Math	254	Apache v2.0	(Zhong et al., 2023)
Hellaswag Val	Hellaswag	Common Sense Reasoning	10,042	MIT	(Zellers et al., 2019)
LogiQA (English)	AGI Eval	Logic	651	GitHub	(Liu et al., 2020)
LSAT-AR	AGI Eval	Law (Analytic Reasoning)	230	MIT	(Wang et al., 2021)
LSAT-LR	AGI Eval	Law (Logical Reasoning)	510	MIT	(Wang et al., 2021)
LSAT-RC	AGI Eval	Law (Reading Comprehension)	260	MIT	(Wang et al., 2021)
MedMCQA Valid	MedMCQA	Medicine	6,150	MIT	(Pal et al., 2022)
SAT-English	AGI Eval	English	206	MIT	(Zhong et al., 2023)
SAT-Math	AGI Eval	Math	220	MIT	(Zhong et al., 2023)

Table 2: Problem sets used to create the multi-domain MCQA exam.

Note: The GitHub repository for LogiQA does not include a license file. However, both the paper and readme.md file state that "The dataset is freely available."

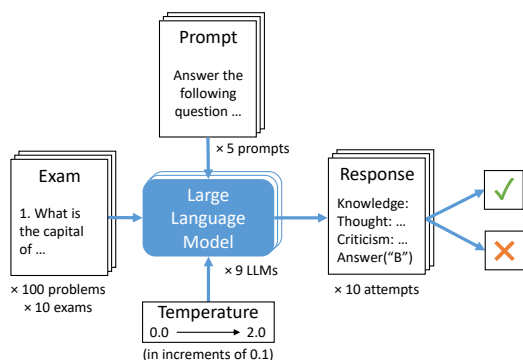


Figure 1: Diagram of the full experiment.

information as efficiently as possible.

First, we instructed GPT-3.5 to complete the 100-question (small) exam using the CoT prompt with temperatures ranging from 0.0 to 2.0 in increments of 0.1. This allowed us to determine the range of viable sampling temperatures to explore.<sup>5</sup>

Performance began to drop rapidly after a temperature of 1.0 until the generated text became incoherent at 1.6. As a result, we stopped the initial temperature sweep at 1.6 and limited the rest of our sweeps from 0.0 to 1.0.

Next, we instructed the other eight LLMs to complete the 100-question (small) exam using the CoT prompt with temperatures from 0.0 to 1.0. This allowed us to determine if the results generalize to other LLMs.

Then, we instructed GPT-3.5 to complete the 100-question (small) exam using each of the five prompts over temperatures from 0.0 to 1.0. This

<sup>5</sup>For this experiment, we fixed all other sampling parameters (e.g., top-k, top-p, etc.) to isolate the effects of temperature.

### Algorithm 1 Full LLM Temperature Experiment

```

1: for each model  $m$  in  $M$  do ▷ 10 models
2:   for each prompt  $p$  in  $P$  do ▷ 5 prompts
3:     for each exam  $e$  in  $E$  do ▷ 10 exams
4:       for each temperature  $\tau$  in  $T$  do ▷ 16 temperatures
5:         for each problem  $q$  in  $Q$  do ▷ 100 problems
6:           for each attempt  $a$  in  $A$  do ▷ 10 attempts
7:             Create the prompt
8:             Answer the question
9:             Record the answer
10:          end for
11:        end for
12:      end for
13:    Save the results
14:  end for
15: end for
16: end for
17: Process the results
18: Analyze the results

```

allowed us to determine if the results generalize over various prompt-engineering techniques.

Finally, we instructed GPT-3.5 to complete the 1,000-question (large) exam using the CoT prompt with temperatures from 0.0 to 1.0. This allowed us to determine if the results were domain-specific or generalized across problem domains.

### 2.4 Metrics

To test our hypothesis, we measured the LLM's correct-answer accuracy as our primary performance metric. For each combination of model, prompt, exam, and temperature, we calculated the accuracy as the number of correct answers from ten attempts at each problem. Then, we computed the average (mean) accuracy across all problems.

To further support our findings, we also measured the similarity of the LLM's responses using a series of text-similarity metrics. These metrics are defined as follows:

- **Jaccard similarity** – the ratio of the intersection to the union of word sets in the output text (Jaccard, 1912).
- **Bag-of-Words (BoW) similarity** – comparison of the frequency of each word, ignoring order (Harris, 1954).
- **TF-IDF similarity** – comparison of word frequency weighted by inverse document frequency (Jones, 1972).
- **Levenshtein similarity** – the number of edits needed to change one string of text into the other (Levenshtein, 1966).
- **BLEU score** – comparison of similarity based on n-gram overlap (Papineni et al., 2001).
- **SBERT similarity** – semantic similarity computed using Sentence-BERT embeddings (Reimers and Gurevych, 2019).

## 2.5 Analysis

We used the Kruskal-Wallis test to evaluate the statistical significance of any changes in accuracy as a function of temperature (Kruskal and Wallis, 1952). We chose the Kruskal-Wallis test because the data (i.e., correct-answer accuracy by question) were not normally distributed. Rather, they were bimodally distributed with centers at 0.0 and 1.0.

## 3 Results

### 3.1 Accuracy vs. Temperature

Our analysis revealed that the problem-solving performance of LLMs remained relatively stable across sampling temperatures from 0.0 to 1.0 for all LLMs, prompt-engineering techniques, and problem domains. Using GPT-3.5 with a CoT prompt on the 1,000-question exam from 0.0 to 1.0, the Kruskal-Wallis test yielded  $H(10) = 10.439$ ,  $p = 0.403$ .

First, we analyzed the performance of GPT-3.5 using the CoT prompt on the 100-question exam. Accuracy remained stable over temperatures from 0.0 to 1.0. However, after a temperature of 1.0, the text rapidly became incoherent, and the accuracy began to drop until it reached zero around a temperature of 1.6. See Figure 2.

Second, we analyzed the performance of all nine LLMs using the CoT prompt on the 100-question exam. Accuracy also remained stable across all of the LLMs, except for Llama 2 7B. The performance




Figure 2: Accuracy by temperature from 0.0 to 1.6 for GPT-3.5 using the CoT prompt on the 100-question exam.

of most LLMs showed a gradual (non-significant) decrease in performance as a function of temperature. See Figure 3 and Table 3.




Figure 3: Accuracy by temperature and model using the CoT prompt on the 100-question exam.

Model	H(10)	p-value
Claude 3 Opus	1.735	0.998
Command R+	1.771	0.998
Gemini Pro 1.0	7.379	0.689
Gemini Pro 1.5	3.119	0.978
GPT-3.5 Turbo	2.042	0.996
GPT-4	3.789	0.956
Llama 2 70B	3.677	0.961
Llama 2 7B	17.086	0.072
Mistral Large	3.069	0.980

Table 3: Kruskal-Wallis test results by model using the CoT prompt on the 100-question exam.

Llama 2 7B did not perform better than statistically random guesses. Its poor performance

was due to generating incorrectly formatted answers (39%) and correctly formatted but incorrect answers (36%). Its all-or-nothing behavior at a temperature of 0.0 versus more random behavior from 0.1 to 1.0 led to a much lower, yet still non-significant, p-value.

Third, we analyzed the performance of GPT-3.5 using each of the five prompts on the 100-question exam. Accuracy remained stable for all temperatures across all prompt-engineering techniques. The CoT prompt outperformed the other four prompts. As a result, we used the CoT prompt for all single-prompt experiments. See Figure 4 and Table 4.

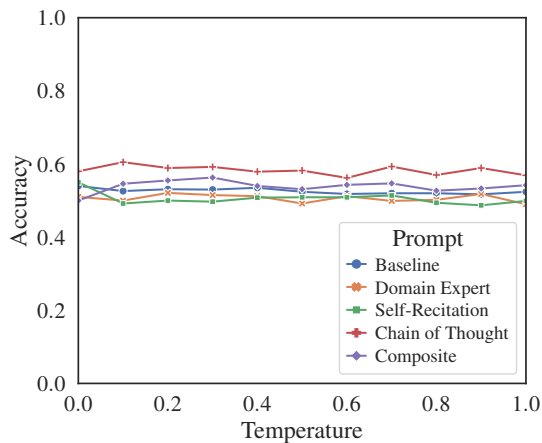


Figure 4: Accuracy by temperature and prompt for GPT-3.5 on the 100-question exam.

Prompt	H(10)	p-value
Baseline	0.420	1.000
Domain Expert	0.548	1.000
Self-recitation	1.403	0.999
Chain of Thought	2.042	0.996
Composite	1.000	1.000

Table 4: Kruskal-Wallis test results by prompt for GPT-3.5 on the 100-question exam.

Finally, we analyzed the performance of GPT-3.5 using the CoT prompt on all ten exams. Accuracy remained stable for all temperatures across all problem domains based on visual analysis. However, the LSAT-AR and SAT-Math exams showed statistically significant differences in the Kruskal-Wallis p-values.<sup>6</sup> See Figure 5 and Table 5.

<sup>6</sup>We considered the ARC Challenge results to be non-significant since they were greater than the significance threshold of 0.05.

We performed the Dunn-Bonferroni test on the LSAT-AR and SAT-Math results (Dunn, 1964). It revealed that the all-or-nothing behavior of responses generated at a temperature of 0.0 versus the more random responses from 0.1 to 1.0 caused the anomaly. The correct-answer accuracy for *each individual problem* varied significantly when evaluated pairwise across temperatures. However, the average accuracy for *all* problems remained similar across temperatures.




Figure 5: Accuracy by temperature and exam for GPT-3.5 using the CoT prompt.

Exam	H(10)	p-value
AQUA-RAT	10.320	0.413
ARC Challenge	16.390	0.089
Hellaswag	4.473	0.924
LogiQA	3.208	0.976
LSAT-AR	37.874	< 0.001
LSAT-LR	7.816	0.647
LSAT-RC	4.037	0.946
MedMCQA	2.334	0.993
SAT-English	3.937	0.950
SAT-Math	21.276	0.019

Table 5: Kruskal-Wallis test results by exam for GPT-3.5 using the CoT prompt.

### 3.2 Text Variability vs. Temperature

To further support our results, we analyzed text variability as a function of temperature. Our findings show a clear trend of decreasing text similarity (thus increasing text variability) as temperature increases. Text similarity decreases rapidly after a temperature of 1.0, corresponding to the rapid decrease in accuracy observed above  $\tau = 1.0$ . See Figure 6.




Figure 6: Text similarity by temperature and metric for GPT-3.5 using CoT prompting on the 100-question exam over sampling temperatures from 0.0 to 1.6.

These results are consistent with our understanding of sampling temperature, indicating that higher temperatures produce more widely varied outputs. Furthermore, these results hold regardless of the LLM, prompt-engineering technique, or problem domain. See Figures 7, 8, and 9.




Figure 7: TF-IDF text similarity by temperature and model using the CoT prompt on the 100-question exam over sampling temperatures from 0.0 to 1.0

## 4 Discussion

### 4.1 Interpretation

Based on these results, changes in temperature from 0.0 to 1.0 do not have a statistically significant effect on the problem-solving performance of LLMs. These results appear to generalize across LLMs, prompt-engineering techniques, and problem domains. However, there are a few potential exceptions to these general findings.




Figure 8: TF-IDF text similarity by temperature and prompt for GPT-3.5 on the 100-question exam over sampling temperatures from 0.0 to 1.0.




Figure 9: TF-IDF text similarity by temperature and exam for GPT-3.5 using the CoT prompt over sampling temperatures from 0.0 to 1.0

Therefore, we recommend setting an LLM's sampling temperature to 0.0 for problem-solving tasks. This temperature maximizes reproducibility without compromising accuracy. In addition, it helps avoid the inevitable drop-off in performance that occurs beyond a temperature of 1.0. However, exceptions to this guidance should be taken into consideration.

### 4.2 Limitations

There were several limitations to our research study due to cost and runtime considerations:

First, our study was limited to a subset of popular LLMs. As a result, our findings may not hold for other LLMs that were excluded from our experiment.

Second, we only explored a subset of prompt-engineering techniques using a single prompt-and-response cycle with one-shot in-context learning. As a result, there may be more complex prompts or agent architectures that better leverage sampling temperature for creativity in their problem-solving capabilities.

Third, our study was limited to a subset of problems, problem domains, and problem-solving tasks. As a result, our findings may not hold for larger data sets, different problem domains, or other types of problem-solving tasks.

Fourth, due to time and cost constraints, we limited our study to two test sets of 1,000 and 100 randomly selected questions from standard benchmarks. These limited sample sizes may have introduced bias into the results. Utilizing a larger and more diverse test set would enhance the statistical reliability of our findings.

Fifth, we had to limit the sampling temperature range we explored from 0.0 to 1.0 for all combinations of models, prompts, and exams, except for GPT-3.5 using CoT prompting on the 100-question exam. As a result, the temperature hyperparameter of other LLMs may operate differently at temperatures above 1.0.

Sixth, we fixed all other sampling parameters (e.g., top-p, top-k, repetition penalty, etc.) to isolate the effect of sampling temperature. As a result, there may be combinations of sampling parameters that result in different outcomes.

Finally, we could only explore a subset of the various combinations of models, prompts, exams, and temperatures. As a result, other combinations of LLMs, prompt-engineering techniques, and problem domains may exist where temperature plays a more important role in problem-solving performance.

### 4.3 Implications

This research study provides empirical evidence that changes in sampling temperature in the range of 0.0 to 1.0 do not significantly impact the problem-solving capabilities of LLMs on MCQA problems.

Answering this question may save AI engineers significant time and resources evaluating various sampling temperatures for their LLM agents and applications. In addition, it may reduce unproductive debates in the prompt-engineering community regarding the optimal sampling temperatures for various problem-solving tasks.

This research also provides broader insights for AI researchers studying model hallucination and problem-solution state-space search with LLMs. Our results show that increasing LLM temperature up to 1.0 does not cause the LLM to hallucinate in ways that lead to incorrect MCQA solutions. In addition, higher temperatures do not appear to improve MCQA solution-space search in ways that lead to correct solutions more often than lower temperatures.

### 4.4 Future Research

To improve upon this research, we propose the following follow-up experiments:

First, we recommend conducting this experiment with additional LLMs. Other proprietary and open-source LLMs may utilize temperature in ways that benefit their specific models but did not benefit the LLMs we tested.

Second, we recommend expanding beyond MCQA problems to other types of problem-solving tasks whose correct answers are more open-ended. The limited effects of sampling temperature in our experiments may have simply resulted from the constraints imposed by the structure of MCQA problems.

Third, we recommend conducting additional experiments with more MCQA problems and problem domains. We recommend specifically targeting tasks and problem domains that require more creative solutions or lateral “out-of-the-box” thinking.

Fourth, we recommend extending the sampling temperature range until accuracy drops to zero for each LLM, prompt, and exam. However, it should be noted that as the generated text becomes more random, the number of tokens in each response increases significantly, leading to a considerable increase in runtime and cost to explore temperatures above 1.0.

Finally, we recommend a more in-depth error analysis to determine if any sub-types of problems within these problem domains benefit from changes to sampling temperature. It is possible that statistical noise or averaging may have hidden individual problems that were sensitive to changes in sampling temperature.

## 5 Conclusion

This research study empirically investigated the effect of sampling temperature on the problem-solving performance of LLMs across multiple prob-



lem domains.

We demonstrated that changes in sampling temperature from 0.0 to 1.0 do not produce statistically significant differences in problem-solving performance on MCQA problems across multiple LLMs, prompt-engineering techniques, and problem domains.

These results have practical implications for AI engineers using LLMs to develop new AI systems. Additionally, they have theoretical implications for AI researchers studying model hallucination and solution-space search with LLMs.

## References

- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. 1985. [A learning algorithm for Boltzmann machines](#). *Cognitive Science*, 9:147–169.
- Anthropic. 2024a. [The claude 3 model family: Opus, sonnet, haiku](#).
- Anthropic. 2024b. [Introducing the next generation of claude anthropic](#).
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? Try ARC, the AI2 reasoning challenge. *ArXiv*.
- Cohere. 2024a. [Command r+](#).
- Cohere. 2024b. [Model card for c4ai command r+](#).
- Olive Jean Dunn. 1964. [Multiple comparisons using rank sums](#). *Technometrics*, 6(3):241–252.
- Gemini-Team. 2023. Gemini: A family of highly capable multimodal models. *arXiv*.
- Gemini-Team. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#).
- Zellig S. Harris. 1954. [Distributional structure](#). *WORD*, 10:146–162.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv*.
- Siqing Huo, Negar Arabzadeh, and Charles L. A. Clarke. 2023. [Retrieving supporting evidence for generative question answering](#). *arXiv*.
- Paul Jaccard. 1912. [The distribution of flora in the alpine zone](#). *New Phytologist*, 11:37–50.
- Karen Spärck Jones. 1972. [A statistical interpretation of term specificity and its application in retrieval](#). *Journal of Documentation*, 28:11–21.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213.
- William H. Kruskal and W. Allen Wallis. 1952. [Use of ranks in one-criterion variance analysis](#). *Journal of the American Statistical Association*, 47(260):583–621.
- Minhyeok Lee. 2023. [A mathematical investigation of hallucination and creativity in GPT models](#). *Mathematics*, 11:2320.
- V.I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, 10:707–710.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. Logiqa: A challenge dataset for machine reading comprehension with logical reasoning. In *International Joint Conference on Artificial Intelligence*.
- Llama-2-Team. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv*.
- Meta. 2023. [Meta and microsoft introduce the next generation of llama | meta](#).
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, Edouard Grave, Yann LeCun, and Thomas Scialom. 2023. Augmented language models: a survey. *arXiv*.
- Microsoft. 2023. [Completions - learn how to generate or manipulate text](#).
- Mistral-AI-Team. 2024. [Au large | mistral ai | frontier ai in your hands](#).
- OpenAI. 2022a. [Introducing ChatGPT](#).
- OpenAI. 2022b. [Models - openai api](#).
- OpenAI. 2023a. [Api temperature change from 0 to 1 to 0 to 2 in 'playground'](#).
- OpenAI. 2023b. [GPT-4](#).
- OpenAI. 2023c. [GPT-4 technical report](#). *arXiv*.
- OpenAI. 2023d. [OpenAI - API reference](#).
- Ankit Pal, Logesh Kumar Umaphathi, and Malaikanan Sankarasubbu. 2022. MedMCQA: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 248–260. PMLR.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2001. BLEU. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02*, page 311. Association for Computational Linguistics.

- Sundar Pichai and Demis Hassabis. 2023. [Introducing gemini: Google’s most capable ai model yet](#).
- Sundar Pichai and Demis Hassabis. 2024. [Introducing gemini 1.5, google’s next-generation ai model](#).
- Vinay Pursnani, Yusuf Sermet, and Ibrahim Demir. 2023. Performance of ChatGPT on the US fundamentals of engineering exam: Comprehensive assessment of proficiency and potential implications for professional environmental engineering practice. *arXiv*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Jessica Shieh. 2024. [Best practices for prompt engineering with OpenAI API](#).
- Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2023. Recitation-augmented language models. In *The Eleventh International Conference on Learning Representations*.
- Chi Wang, Susan Xueqing Liu, and Ahmed H Awadallah. 2023a. Cost-effective hyperparameter optimization for large language model generation inference.
- Pei-Hsin Wang, Sheng-Iou Hsieh, Shih-Chieh Chang, Yu-Ting Chen, Jia-Yu Pan, Wei Wei, and Da-Chang Juan. 2020. Contextual temperature for language modeling. *arXiv*.
- Rui Wang, Hongru Wang, Fei Mi, Yi Chen, Ruifeng Xu, and Kam-Fai Wong. 2023b. Self-critique prompting with large language models for inductive instructions. *arXiv*.
- Siyuan Wang, Zhongkun Liu, Wanjun Zhong, Ming Zhou, Zhongyu Wei, Zhumin Chen, and Nan Duan. 2021. From lsat: The progress and challenges of complex reasoning. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 30:2201–2216.
- Ian Ward. 2014. [JSON lines](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. *arXiv*.
- Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C. Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with ChatGPT. *arXiv*.
- Frank F Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. [A systematic evaluation of large language models of code](#). In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, pages 1–10. Association for Computing Machinery.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. AGIEval: A human-centric benchmark for evaluating foundation models. *ArXiv*.
- Yuqi Zhu, Jia Li, Ge Li, YunFei Zhao, Jia Li, Zhi Jin, and Hong Mei. 2023. Improving code generation by dynamic temperature sampling. *arXiv*.

## A Appendix

```
[System Prompt]
You are an expert in {{expertise}}.
Your task is to answer the following multiple-choice questions.
First, you should recite all of the relevant knowledge you have about the question and each option.
Next, you should think step-by-step through the problem to ensure you have the correct answer.
Then, you should critically evaluate your thoughts to identify any flaws in your facts, logic, and reasoning.
Finally, you MUST answer the question using the following format 'Action: Answer("[choice]")'
The parameter [choice] is the letter or number of the answer you want to select (e.g. "A", "B", "C", or "D")
For example, 'Answer("C")' will select choice "C" as the best answer.
The answer MUST ALWAYS be one of the available choices; it CANNOT be "None of the Above".
If you think the answer is "none of the above", then you MUST select the most likely answer.

[Example Problem]
Question: What is the capital of the state where Johns Hopkins University is located?
Choices:
A: Baltimore
B: Annapolis
C: Des Moines
D: Las Vegas

[Example Solution]
Knowledge:
Johns Hopkins University is located in Baltimore, Maryland.
A: Baltimore is a city located in the State of Maryland, but it is not the capital of Maryland.
B: Annapolis is the capital of the State of Maryland.
C: Des Moines is a city located in the State of Iowa and is the capital of Iowa.
D: Las Vegas is located in the State of Nevada, but it is not the capital of Nevada.
Thought:
Johns Hopkins University is located in Baltimore.
Baltimore is a city located in the state of Maryland.
The capital of Maryland is Baltimore.
Therefore, the capital of the state where Johns Hopkins University is located is Baltimore.
The answer is A: Baltimore.
Criticism:
You are correct that Johns Hopkins is located in the State of Maryland.
However, the capital of Maryland is Annapolis, not Baltimore.
So, the correct answer is actually B: Annapolis.
Action: Answer("B")
```

Figure 10: Sample of the composite system prompt with a one-shot example (i.e., problem-and-solution pair).

```
{
  "source": "arc/arc-challenge-test",
  "source_id": 1,
  "topic": "Science",
  "context": "",
  "question": "An astronomer observes that a planet rotates faster
    after a meteorite impact. Which is the most likely effect
    of this increase in rotation?",
  "choices": {
    "A": "Planetary density will decrease.",
    "B": "Planetary years will become longer.",
    "C": "Planetary days will become shorter.",
    "D": "Planetary gravity will become stronger." },
  "answer": "C",
  "solution": ""
}
```

Figure 11: Sample of an MCQA problem in JSON-L format – with whitespace added for readability.