

PERSONA SWITCH: Mixing Distinct Perspectives in Decoding Time

Junseok Kim¹ Nayeong Yang¹ Kyomin Jung^{1†}

¹Seoul National University

{kim.junseok, yny0506, kjung}@snu.ac.kr

Abstract

Role-play prompting is known to steer the behavior of language models by injecting a persona into the prompt, improving their zero-shot reasoning capabilities. However, such improvements are inconsistent across different tasks or instances. This inconsistency suggests that zero-shot and role-play prompting may offer complementary strengths rather than one being universally superior. Building on this insight, we propose **Persona Switch**, a novel decoding method that dynamically combines the benefits of both prompting strategies. Our method proceeds step-by-step, selecting the better output between zero-shot and role-play prompting at each step by comparing their output confidence, as measured by the logit gap. Experiments with widely-used LLMs demonstrate that Persona Switch consistently outperforms competitive baselines, achieving up to 5.13% accuracy improvement. Furthermore, we show that output confidence serves as an informative measure for selecting the more reliable output.¹

1 Introduction

Recent works have explored guiding LLM behavior by assigning personas to the model, known as role-play prompting (Shanahan et al., 2023; Han et al., 2022). This line of studies has suggested that role-play prompting improves language model performance, such as zero-shot reasoning capability, showing empirical observations reveal that the model’s responses can vary depending on the presence of persona (Xu et al., 2023; Kong et al., 2024). Variations in response suggest that the presence of a persona influences the internal representation of the LLM, potentially due to patterns implicitly learned from the model’s training process.

However, subsequent observations indicate that role-play prompting does not always guarantee con-

[†]Corresponding author.

¹Our code and data have been released on <https://github.com/junseokkim00/PersonaSwitch>.

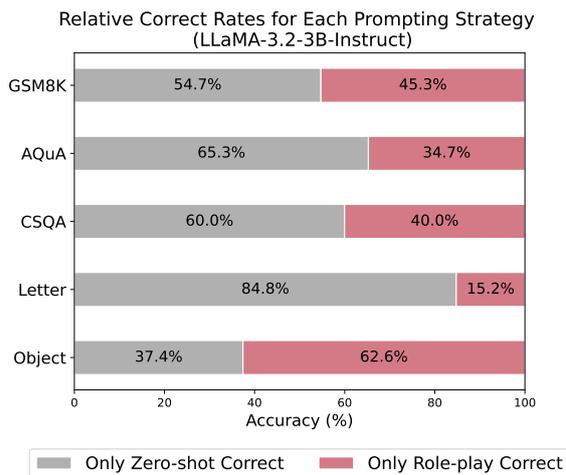


Figure 1: **Relative Correct Rates Show No Clear Superiority Between Prompting Strategies.** Each bar represents the proportion of instances where only one strategy answered correctly, highlighting that neither one demonstrates clear superiority.

sistent benefits. This is because prepending a role-play prompt to the context does not contribute essential information that facilitates problem solving. For instance, consider a simple math problem like “What is $4 + 8$?”. Whether the model receives the question without any persona (zero-shot) or while role-playing as a *math teacher*, the key information given to the model is unchanged, and the task’s complexity is unaffected either. To illustrate this, Figure 1 shows the proportion of cases where only one prompting strategy produces the correct answer. The result reveals that the superiority depends on each instance, with neither strategy consistently outperforming the other. Rather, they possess complementary strengths, underscoring the potential of a unified approach that integrates the strengths of both strategies to achieve better performance.

Motivated by these observations, we propose **Persona Switch**, a novel decoding method that considers both zero-shot and role-play prompting perspectives. Specifically, our method compares

the single-step outputs generated from different prompting strategies and selects the more reliable one based on the logit gap. The final response is incrementally constructed by repeating this step-by-step selection process, effectively mixing perspectives of both prompting strategies during decoding. Evaluations on five benchmarks demonstrate consistent improvement, achieving 5.13% of average accuracy gains over the strongest baseline in widely-used LLM. We also find that the logit gap serves as an informative criterion for selecting reliable outputs, leading to improved performance. Furthermore, incorporating only informative tokens to assess output confidence enables more fine-grained selection of reliable outputs, further contributing to performance gains.

2 Related Works

Role-Play Prompting in LLM Prompting has become a widely used paradigm for adapting language models to various NLP tasks, with zero-shot, few-shot, and instruction-based prompting, showing strong empirical performance (Brown et al., 2020; Wei et al., 2021; Kojima et al., 2022). Recent works have explored role-play prompting, where LLMs follow persona-based instructions to simulate specific behaviors (Han et al., 2022; Shanahan et al., 2023; Xu et al., 2023; Kong et al., 2024). These approaches report performance gains across specific tasks, though such benefits are not always consistent or well understood (Zheng et al., 2024; Kim et al., 2025). Recent work also highlights that role-play prompting can introduce biases, raising concerns about its robustness and consistent effectiveness across tasks (Salewski et al., 2023; Gupta et al., 2023). Consequently, the inconsistent gains from role-play prompting suggest that it may complement non-role-play strategies, motivating a unified approach that leverages the strengths of both strategies.

3 Methods

3.1 Problem Definition

Role-play prompting is known to effectively steer the model’s behavior by inserting personas into the prompt. However, while adding a persona can change how the model responds, it does not modify the key information available to solve the problem. We define this invariance as the principle of *Persona-Agnostic Semantic Invariance*.

Persona-Agnostic Semantic Invariance. *Given the same input, introducing a persona does not alter the semantic content of the prompt, thus does not add or remove any task-relevant information.*

Therefore, if zero-shot and role-play prompting yield different outputs, there is no basis for claiming that either one is inherently superior. This lack of clear superiority introduces ambiguity when choosing between outputs. To address this, we propose **Persona Switch**, a decoding method that dynamically selects the better output between zero-shot and role-play prompting at each step, based on the model’s internal confidence signal (logit gap).

3.2 Persona Switch

This section introduces how Persona Switch is executed. Specifically, Persona Switch consists of two steps: **Output Candidate Generation** and **Output Confidence Comparison**.

Output Candidate Generation. Given an input x , a language model \mathcal{M} , and a reasoning trace \hat{y}^t (initially an empty string) at step t , we generate an output using greedy decoding from zero-shot y_{base} and role-play prompting y_{persona} . We adopt the role-play prompt p introduced by Kong et al. (2024) (dataset-specific prompts in Appendix A). Following Cao et al. (2024), we use $\backslash n \backslash n$ as a stop string to enforce single-step generation. Implementation details about single-step generation are in Appendix B. The output y_m ($m \in \{\text{base}, \text{persona}\}$) for each prompting strategy is defined as follows:

$$\begin{aligned} y_{\text{base}} &= \mathcal{M}(x; \hat{y}^t) \\ y_{\text{persona}} &= \mathcal{M}(p; x; \hat{y}^t) \end{aligned} \quad (1)$$

Output Confidence Comparison. This section describes how to select the final output y_{final}^t for step t . Let y_m denote a sequence of tokens $(y_{m,1}, y_{m,2}, \dots, y_{m,n})$, where n is the length of y_m . We compare their output confidence to select the better one between the two outputs. This approach is motivated by the observation that low-confidence generations are often correlated with incorrect or hallucinated outputs, as suggested by prior studies that analyzed model confidence estimation and generation quality (Kadavath et al., 2022; Zhang et al., 2023; Wang and Zhou, 2024). Based on this insight, we adopt the logit gap, a token-level confidence measure that computes the difference between the highest and second-highest probabilities for each token. Formally, the logit gap

Model	Method	GSM8K	AQuA	CSQA	Letter	Object	Average
Gemma-2 (2B)	Greedy	63.08	<u>31.89</u>	<u>67.24</u>	<u>11.00</u>	31.87	<u>41.01</u>
	Top-p	61.18	28.74	66.18	6.00	29.07	38.23
	Top-k	59.59	28.74	66.01	4.80	30.13	37.86
	Multinomial	59.21	31.10	65.77	5.20	<u>32.27</u>	38.71
	Role-Play Prompting	<u>63.61</u>	28.35	53.15	6.20	28.4	35.94
	Low-gap Selection	61.86	31.49	57.33	5.40	31.07	37.43
	Random Selection	63.25	37.93	61.73	7.86	32.44	40.64
	PERSONA SWITCH	64.14	40.16	67.98	13.60	34.67	44.11 (+3.10)
LLaMA-3.2 (3B)	Greedy	<u>73.09</u>	40.94	69.45	65.40	53.47	<u>60.47</u>
	Top-p	71.72	37.40	65.36	61.00	51.07	57.31
	Top-k	67.55	<u>41.34</u>	63.31	54.80	46.53	54.71
	Multinomial	73.01	38.58	64.78	<u>66.80</u>	51.33	58.90
	Role-Play Prompting	71.49	29.53	66.99	56.00	<u>64.13</u>	57.63
	Low-gap Selection	72.71	35.83	67.49	55.20	58.67	57.98
	Random Selection	74.73	40.94	68.28	62.53	61.15	61.52
	PERSONA SWITCH	76.65	44.88	<u>69.12</u>	72.80	64.53	65.60 (+5.13)
LLaMA-3.1 (8B)	Greedy	<u>83.62</u>	51.97	71.42	<u>82.80</u>	58.80	69.72
	Top-p	79.91	54.72	68.71	72.20	64.93	68.09
	Top-k	79.83	51.97	68.63	66.80	61.33	65.71
	Multinomial	83.17	<u>55.12</u>	69.62	78.00	64.67	<u>70.12</u>
	Role-Play Prompting	77.03	42.13	<u>73.05</u>	74.20	71.33	67.55
	Low-gap Selection	83.70	59.06	70.93	75.40	58.13	69.44
	Random Selection	84.71	58.40	72.02	77.47	64.67	71.45
	PERSONA SWITCH	85.75	59.06	73.30	84.20	<u>70.40</u>	74.54 (+4.42)

Table 1: **Experimental results on five reasoning datasets.** We report accuracy (%) for each method. Bolded numbers represent the best performance for each dataset, while underlined numbers denote the second-best performance. Parentheses indicate Persona Switch’s absolute improvement over the best baseline for each model.

$\Delta y_{m,j}$ for each token $y_{m,j}$ is computed as follows:

$$\Delta y_{m,j} = P(y_{m,j}^1 | x, \hat{y}^t, y_{m,<j}) - P(y_{m,j}^2 | x, \hat{y}^t, y_{m,<j}) \quad (2)$$

where $y_{m,j}^1$ and $y_{m,j}^2$ indicate the top two tokens at the j -th decoding step. By summing these logit gaps, we can compute the average logit gap Δy_m :

$$\Delta y_m = \frac{1}{n} \sum_{j=1}^n \Delta y_{m,j} \quad (3)$$

After computing Δy_m for each prompting strategy, we select the output with the larger value as the final output y_{final}^t for step t . The reasoning trace is then updated as $\hat{y}^{t+1} = \hat{y}^t \parallel y_{\text{final}}^t$ via concatenation. This process is repeated iteratively until y_{final}^t is an *eos* token, after which the final answer is extracted from the resulting reasoning trace.

$$y_{\text{final}}^t = \begin{cases} y_{\text{base}} & \text{if } \Delta y_{\text{base}} > \Delta y_{\text{persona}} \\ y_{\text{persona}} & \text{otherwise} \end{cases} \quad (4)$$

Answer Extraction. Identifying the exact answer span is challenging. We follow Kojima et al. (2022) by concatenating an answer trigger prompt to elicit only the final answer. The answer trigger prompts for each dataset are in Appendix C.

4 Experiments

4.1 Experimental setup

Experimental Settings. We evaluate our method on five widely-used reasoning benchmarks: GSM8K (Cobbe et al., 2021), AQuA-RAT (Ling et al., 2017), CSQA (Talmor et al., 2019), Last Letter Concatenation (Wei et al., 2022) and Tracking Shuffled-Objects from BIG-bench (Srivastava et al., 2022)). We use two instruction-tuned variants of the LLaMA-3 family, namely LLaMA-3.2-3B-Instruct and LLaMA-3.1-8B-Instruct, and one from Gemma-2 family, Gemma-2-2B-it, following recent studies demonstrating competitive zero-shot reasoning performance (Grattafiori et al., 2024; Team et al., 2024). We compare our method against five baselines (Zero-Shot with Greedy, Top-p, Top-k, Multinomial sampling, and Role-Play Prompting (Kong et al., 2024)), as well as two Persona Switch variants: Low-gap Selection (selects the output with lower average logit gap), and Random Selection (randomly selects between two outputs). Implementation details are in Appendix D.

Dataset	Question	Zero-Shot Response	Role-Play Response	PERSONA SWITCH
GSM8K	A DVD can be played 1000 times . . . How many total times can both DVDs be played before they break?	To find the total number of times both DVDs . . . So, both DVDs can be played a total of 2000 times before they break. ✗	To find the total number of times both DVDs . . . So, both DVDs can be played a total of 1509 times before they break. ✓	To find the total number . . . The first DVD can . . . The second DVD can . . . Therefore, both DVDs can be played a total of 1509 times before they break. ✓
Letter	Take the last letters of each words in "Rich Parker Edward Milagros" and concatenate them.	To find the last letters of each word, . . . Now, let's concatenate these letters: H-R-D-S . ✓	To find the last letters of each word, . . . The resulting word is " Hards ". ✗	To find the last letters of each word, . . . Now, let's concatenate these letters: H-R-D-S . ✓
Object	Alice, Bob, and Claire are holding a white . . . Which choice is true? Answer Choices: (A) black ball. (B) green present. (C) orange ball.	Let's break down the gift swaps: 1. Alice and Bob swap gifts: . . . The correct answer is (A) black ball . ✗	Let's follow the gift exchange process: 1. Alice and Bob swap gifts: . . . The correct answer is (A) black ball . ✗	Let's follow the gift exchange process: 1. Alice and Bob swap gifts: . . . The correct answer is (B) green present . ✓

Table 2: **Qualitative Analysis on examples of GSM8K, Letter, and Object datasets.** Each example includes responses from zero-shot, role-play prompting, and Persona Switch. **Blue** indicates response from zero-shot, and **Red** indicates response from role-play prompting. ✓ and ✗ indicate correct and wrong, respectively.

4.2 Results and analysis

Main Results. As shown in Table 1, Persona Switch consistently outperforms competitive baselines. Our method achieves the highest overall performance regarding average accuracy, surpassing the strongest baseline by 5.13% and 4.42%, respectively. Furthermore, compared to both low-gap selection and random selection, Persona Switch demonstrates superior accuracy across all datasets, underscoring the efficacy of our logit-based answer selection mechanism. These results validate that selecting responses based on the averaged logit gap yields robust and reliable improvements in multi-domain reasoning tasks.

Analysis of Output Granularity. Choosing a proper output unit for our method is crucial, as the final response can vary depending on how often outputs of zero-shot and role-play prompting are compared. To assess the impact of output granularity, we evaluated our method at the (1) token, (2) step, and (3) full-sequence level using LLaMA-3.2-3B-Instruct. Table 3 shows that step-level selection performs best, indicating that evaluating intermediate reasoning steps allows a more diverse integration of zero-shot and role-play perspectives compared to full-sequence selection, leading to improved performance. On the token level, its evaluation tends to underperform as it requires selecting the correct option for every token, where individual tokens lack sufficient context to distinguish the correct reasoning path. Additionally, we also compared confidence using maximum logit, and the consistently better performance of logit gap confirms it as a more informative measure.

Analysis of Filtering Stopwords. Our approach computes the average logit gap across all tokens when comparing two outputs. However, not all

Metric	Granularity	GSM8K	AQuA	Letters
Max logit	Step	76.27	42.13	70.60
Logit gap	Token	75.13	41.73	71.40
	Full-Seq	74.83	38.58	63.40
	Step (ours)	76.65	44.88	72.80

Table 3: **Performance across granularities and metrics on GSM8K, AQuA, and Letters using LLaMA-3.2-3B-Instruct.** Bold indicates the best performance for each benchmark.

tokens contribute equally to output semantics, as some carry essential information while others, such as stopwords, do not. To assess the impact of removing uninformative tokens, we conduct experiments using two filtering strategies: (1) **stopwords**, excluding stopwords based on NLTK’s predefined list², and (2) **only numbers**, retaining only numerical values. Table 4 shows that applying stopword filtering yields a slight performance gain, suggesting that irrelevant tokens can obscure confidence signals. In contrast, using only numerical tokens degrades performance, likely due to the loss of linguistically meaningful context. Findings suggest that the types of meaningful tokens vary across benchmarks, and applying a task-specific filtering strategy can potentially enhance performance.

Model Size	Filtered wordlist	GSM8K	AQuA	Object
3B	N/A	76.65	44.88	64.53
	only numbers	76.12	40.94	-
	stopwords	76.27	46.06	64.80
8B	N/A	85.75	59.06	70.40
	only numbers	85.52	59.84	-
	stopwords	85.90	60.24	70.93

Table 4: **Result for different filtered wordlists on GSM8K, AQuA, and Object datasets.** Bold indicates the best performance per column.

²<https://www.nltk.org/howto/corpus.html>

Qualitative Analysis. Table 2 presents several examples from three benchmarks using LLaMA-3.1-8B-Instruct. In the GSM8K and Last Letter examples, we observe that Persona Switch can produce a better answer when one of the two (zero-shot or role-play prompting) is correct. This demonstrates its ability to identify and favor the more reliable reasoning path. Furthermore, in the Object Tracking example, even when both zero-shot and role-play responses are incorrect, the Persona Switch successfully produces the correct answer by integrating reasoning steps from both outputs. These cases illustrate that the Persona Switch enhances model performance by selectively adopting the better step between the two. Full results for each example can be found in Appendix H.

5 Conclusion

We present Persona Switch, a dynamic decoding strategy that leverages the strengths of both zero-shot and role-play prompting by selecting the better output at each step based on the logit gap. Notably, it can be easily implemented without extra training and shows performance gains compared to using either prompting strategy alone. Ultimately, Persona Switch contributes a simple yet effective decoding-time mechanism, offering a practical framework for integrating distinct prompting strategies.

Limitations

While our method demonstrates promising results, certain limitations suggest potential directions for future work. In our experiments, we adopted fixed role-play prompts from prior studies and used their corresponding datasets, which have a clearly defined ground-truth answer. This setting allowed us to assume persona-agnostic semantic invariance, where the underlying essential information of each input remains consistent regardless of the persona used. However, in more open-ended or subjective tasks where ground-truth answers are inherently ambiguous, the effectiveness of persona-agnostic approaches may vary, motivating further exploration of a framework suited to such settings. Furthermore, our evaluation is limited to zero-shot reasoning tasks. Although our logit-based stepwise selection framework shows strong performance in this context, it remains an open question whether the same approach would be effective in other settings, such as domains beyond reasoning.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No.RS-2023-00229780, Development of Artificial Intelligence Technology for Process-focused Evaluation(Student’s Learning Diagnosis)). K. Jung is with ASRI, Seoul National University, Korea.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Lang Cao, Chao Peng, Renhong Chen, Wu Ning, Ying-tian Zou, and Yitong Li. 2024. Step guided reasoning: Improving mathematical reasoning using guidance generation and step reasoning. *arXiv preprint arXiv:2410.19817*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2023. Bias runs deep: Implicit reasoning biases in persona-assigned llms. *arXiv preprint arXiv:2311.04892*.
- Seungju Han, Beomsu Kim, Jin Yong Yoo, Seokjun Seo, Sangbum Kim, Enkhbayar Erdenee, and Buru Chang. 2022. Meet your favorite character: Open-domain chatbot mimicking fictional characters with only a few utterances. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5114–5132.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and 1 others. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Junseok Kim, Nakyeong Yang, and Kyomin Jung. 2025. Persona is a double-edged sword: Rethinking the

- impact of role-play prompts in zero-shot reasoning tasks. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 848–862.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang, and Xiaohang Dong. 2024. Better zero-shot reasoning with role-play prompting. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4099–4113.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167.
- Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2023. In-context impersonation reveals large language models’ strengths and biases. *Advances in neural information processing systems*, 36:72044–72057.
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role play with large language models. *Nature*, 623(7987):493–498.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adria Garriga-Alonso, and 1 others. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Xuezhi Wang and Denny Zhou. 2024. Chain-of-thought reasoning without prompting. *arXiv preprint arXiv:2402.10200*.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou, Yongdong Zhang, and Zhendong Mao. 2023. Expertprompting: Instructing large language models to be distinguished experts. *arXiv preprint arXiv:2305.14688*.
- Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou, Xinbing Wang, and Luoyi Fu. 2023. Enhancing uncertainty-based hallucination detection with stronger focus. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 915–932.
- Mingqian Zheng, Jiaxin Pei, Lajanugen Logeswaran, Moontae Lee, and David Jurgens. 2024. When” a helpful assistant” is not really helpful: Personas in system prompts do not improve performances of large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15126–15154.

A Dataset-Specific Role-Play Prompts

In our role-play prompting setup, we adopt dataset-specific prompts that were previously designed and validated for task effectiveness by Kong et al. (2024), as shown in Table 5. Each prompt consists of two messages: 1) **role-setting prompt**, which assigns a specific persona to the LLM through the user message, and 2) **role-feedback prompt**, which prompts the model to acknowledge its assigned persona by reinforcing the role through the assistant message.

B Single Step Output Generation

To help understand our approach to generating the output as a single step from Output Candidate Generation, we include a Python code example of using stopping criteria and concatenating reasoning trace \hat{y} when inferring. More details can be found in the API documentation³ of HuggingFace.

```
# A code example of generating output
reasoning_trace = previous outputs

# 1. Apply chat template to the conversation
conversation = [
    {"role": "user", "content": question},
    {"role": "assistant", "content": reasoning_trace}
]
tokenized_chat = tokenizer.apply_chat_template(
    conversation,
    tokenize=False,
    add_generation_prompt=False,
    return_tensors="pt"
)

# 2. Remove the end-of-turn token and encode it
tokenized_chat = tokenized_chat[:-len("<|eot_id|>")]
model_inputs = tokenizer(
    [tokenized_chat],
    return_tensors="pt"
)

# 3. Define stopping criteria
stop_criteria = StoppingCriteria(
    tokenizer,
    model_inputs['input_ids'].shape(1),
    r"\n\n"
)
stoptlist = StoppingCriteriaList([stop_criteria])

# 4. Generate the output with greedy decoding setting
output = model.generate(
    **model_inputs,
    max_new_tokens=256,
    stopping_criteria=stoptlist,
    do_sample=False
)
```

C Answer Extraction

We extract the final answer by following the method of Kojima et al. (2022). Specifically, we concatenate the output with the dataset-specific answer trigger prompt, triggering the model to only

³https://huggingface.co/docs/transformers/internal/generation_utils

return the answer. Answer trigger prompts for each dataset could be found in Table 6.

D Baselines and Variants

In this section, we introduce the specific configuration for each baseline used in the main result. We ran the method three times with different seeds (1,2, and 3) and report the average performance.

Greedy decoding : The model selects the token with the highest probability at each generation. This method is deterministic and often yields fluent but potentially repetitive or overly conservative outputs, as it does not explore alternative candidates.

Top- p sampling (nucleus sampling) : This method samples the token from the smallest possible set of tokens whose cumulative probability exceeds a predefined threshold p . We set $p = 0.9$, meaning the model considers a dynamic set of high-probability tokens at each step, encouraging more diverse outputs while maintaining fluency.

Top- k sampling : The model restricts sampling to the top k most probable tokens at each step, and selects one based on their normalized probabilities. We used $k = 10$ in our experiments. This method controls the randomness of generation by limiting the candidate pool, leading to outputs that are both diverse and locally coherent.

Multinomial sampling (temperature sampling) : This method samples from the full distribution, controlled by a temperature parameter τ . A lower temperature makes the distribution sharper, favoring high-probability tokens, while a higher temperature leads to more randomness. We used $\tau = 0.6$, which is the model's default value.

Low-gap selection : As a variant of our method, this approach selects outputs based on the average logit gap between zero-shot and role-play prompting outputs for each step. Specifically, it computes the logit difference for each token and selects the output with the lower average gap.

Random selection : Another variant of our method, this approach randomly selects between the zero-shot and role-play prompting outputs at each step to construct the final response. This method introduces stochasticity and serves as a baseline for comparison with our method.

Task	Prompt
GSM8K	User Message: From now on, you are an excellent teacher and are teaching your students to get a new word by concatenating the last letters of several words. I am one of your students and want to ask you a related question.
	Assistant Message: Of course! I'd be happy to help you with any questions you have about creating new words by concatenating the last letters of several words. Please go ahead and ask your question, and I'll do my best to assist you.
AQuA	User Message: From now on, you are an excellent teacher and are teaching your students to get a new word by concatenating the last letters of several words. I am one of your students and want to ask you a related question.
	Assistant Message: Of course! I'd be happy to help you with any questions you have about creating new words by concatenating the last letters of several words. Please go ahead and ask your question, and I'll do my best to assist you.
CSQA	User Message: From now on, you are a contestant in the general knowledge quiz contest and always answer all kinds of common sense questions accurately. I am the moderator of the game and the final is about to start.
	Assistant Message: That sounds like an exciting challenge! I'm ready to participate in the quiz contest as a contestant. Please go ahead and start the final round—I'm here to provide accurate answers to your common sense questions.
Letter	User Message: From now on, you are an excellent teacher and are teaching your students to get a new word by concatenating the last letters of several words. I am one of your students and want to ask you a related question.
	Assistant Message: Of course! I'd be happy to help you with any questions you have about creating new words by concatenating the last letters of several words. Please go ahead and ask your question, and I'll do my best to assist you.
Object	User Message: From now on, you are a recorder. Alice, Bob, and Claire invite you to record a game. They will exchange their stuff in order, and you (the recorder) will fully record the whole process and tell them what they end up with.
	Assistant Message: Certainly! I will act as a recorder and document the game in which Alice, Bob, and Claire will exchange their items. Please provide me with the specific order in which they will exchange their belongings, and I will keep track of the process and inform you of what each person ends up with at the end.

Table 5: Prompts for GSM8K, AQuA, Commonsense Question Answering, Last Letter Concatenation, and Tracking Shuffled Objects. For each task, the upper cell contains the role-setting prompt and the lower cell presents the role-feedback prompt.

Datasets	Answer Trigger
GSM8K	Therefore, the answer (arabic numerals) is
AQuA, CSQA	Therefore, among A through E, the answer is
Object	Therefore, among A through C, the answer is
Letter	Therefore, the final answer is

Table 6: Answer trigger prompts for each dataset.

E Dataset Details

In this section, we introduce four different domains, including five distinct datasets used for our experiments. Specific details are shown in Table 7.

Mathematical Reasoning We use **GSM8K** (Cobbe et al., 2021), providing math word problems and **AQuA-RAT** (Ling et al., 2017) for multiple choice questions. All questions in this dataset require mathematical reasoning knowledge.

Commonsense Reasoning We leverage **CSQA** (Talmor et al., 2019), providing multiple choice questions. This dataset requires reasoning based on commonsense knowledge to solve the problem.

Symbolic Reasoning We use **Last Letter Concatenation** (Wei et al., 2022), which requires concatenating the last letter of the given words in order.

Others We utilize **Tracked Shuffling Objects** from Big-Bench (Srivastava et al., 2022), which provides a sequence of trading history of the object, and ask for the object's final position.

F Impact of Role-Play Prompt design

We conducted additional experiments to evaluate the performance of Persona Switch using multiple variants of the role-play prompt template (Prompts 1, 2, and 3), which differ in verbosity and structure, following Kong et al. (2024). Specifically, Prompt 1 consists solely of the user message, Prompt 2

Dataset	Answer Format	N_q	L_q	License
GSM8K	arabic number	1319	46.9	MIT License
AQUA-RAT	option (A-E)	254	51.9	Apache-2.0
CommonsenseQA	option (A-E)	1221	27.8	Unspecified
Last Letters	string	500	15.0	-
Shuffled Objects	option (A-C)	750	91.1	Apache-2.0

Table 7: Relevant information of five datasets. N_q denotes the number of questions in each dataset. L_q denotes the average words of questions in each dataset.

Dataset	Prompt Type	Prompt	Accuracy
AQuA	Prompt 1	User Message: From now on, you are an excellent math teacher and always teach your students math problems correctly. And I am one of your students.	42.91 %
	Prompt 2	User Message: From now on, you are an excellent math teacher and always teach your students math problems correctly. And I am one of your students.	42.91 %
		Assistant Message: That’s great to hear! As your math teacher, I’ll do my best to explain mathematical concepts correctly so that you can understand them easily. Feel free to ask any math problems or questions you have, and I’ll be glad to assist you.	
Prompt 3 (Ours)	User Message: From now on, you are an excellent math teacher and always teach your students math problems correctly. And I am one of your students.	Assistant Message: That’s great to hear! As your math teacher, I’ll do my best to explain mathematical concepts correctly so that you can understand them easily. Feel free to ask any math problems or questions you have, and I’ll be glad to assist you. Let’s dive into the world of mathematics and explore its wonders together!	44.88 %
	Assistant Message: Of course! Please go ahead and ask your question, and I’ll do my best to assist you.		
Letter	Prompt 1	User Message: From now on, you are an excellent teacher and are teaching your students to get a new word by concatenating the last letters of several words. I am one of your students and want to ask you a related question.	69.8 %
	Prompt 2	User Message: From now on, you are an excellent teacher and are teaching your students to get a new word by concatenating the last letters of several words. I am one of your students and want to ask you a related question.	71.2 %
		Assistant Message: Of course! Please go ahead and ask your question, and I’ll do my best to assist you.	
Prompt 3 (Ours)	User Message: From now on, you are an excellent teacher and are teaching your students to get a new word by concatenating the last letters of several words. I am one of your students and want to ask you a related question.	Assistant Message: Of course! I’d be happy to help you with any questions you have about creating new words by concatenating the last letters of several words. Please go ahead and ask your question, and I’ll do my best to assist you.	72.8 %
	Assistant Message: Of course! Please go ahead and ask your question, and I’ll do my best to assist you.		

Table 8: **Variants of persona prompts used in our experiments.** Across AQuA and Last Letter, the more detailed Prompt 3 yields the best accuracy, and is therefore adopted as the default in our main experiments.

includes a user–assistant message pair, and Prompt 3 extends Prompt 2 with an additional assistant sentence, which serves as the default in our main experiments. As presented in Table 8, Prompt 3 resulted in the strongest performance across both AQuA and Last Letter tasks. These results suggest that employing a more detailed role-play prompt was associated with improved performance. Based on these findings, we adopted Prompt 3 as the default configuration across all benchmarks.

G More ablation studies

G.1 Role-Playing vs. Stochastic decoding

To examine whether the performance gains of Persona Switch can be attributed solely to increased decoding diversity, we replace role-playing prompting with standard sampling-based decoding strategies, including top- p , top- k , and multinomial sampling, while keeping the rest of the framework unchanged. As shown in Table 9, sampling-based

Decoding Configuration	AQuA	Letter
Greedy + Top- p ($p = 0.9$)	44.49	69.8
Greedy + Top- k ($k = 10$)	38.58	64.4
Greedy + Multinomial ($\tau = 0.6$)	40.16	66.6
Greedy + Role-Play (PersonaSwitch)	44.88	72.8

Table 9: Ablation study comparing sampling-based decoding strategies with role-playing prompting in Persona Switch using LLaMA-3.2-3B-Instruct.

variants consistently underperform Persona Switch on both AQuA and Letter datasets, despite introducing stochastic diversity. This indicates that Persona Switch benefits not merely from sampling multiple candidates within a single output distribution but from role-playing prompts that induce qualitatively distinct reasoning trajectories by shifting the underlying generation distribution, resulting in more coherent and complementary candidate solutions.

H Full Results of Qualitative Analysis

Table 10, 11, and 12 illustrate the full result for the example GSM8K, Letter, and Object, respectively. Blue and Red colored text indicates zero-shot output and role-play output, respectively.

I Use of AI Tools

During the preparation of this paper, AI tools (e.g., OpenAI’s ChatGPT) were used in a limited, supporting capacity. Specifically, they assisted in enhancing the clarity and fluency of the text and in suggesting relevant keywords during the writing process. All conceptual ideas, experimental designs, implementations, analyses, and final interpretations were developed entirely by the authors. The authors independently verified all cited references, and no citation was included solely based on AI-generated content. No private, unpublished, or sensitive information was shared with AI tools beyond what is explicitly described in this paper.

Question: A DVD can be played 1000 times before it breaks. There are two DVDs in the public library, one has been played 356 times and the other has been played 135 times. How many total times can both DVDs be played before they break?

Zero-Shot: To find the total number of times both DVDs can be played before they break, we need to add the number of times each DVD has already been played to the number of times each can be played before breaking, and then add those two totals together.

DVD 1 has been played 356 times and can be played 1000 times before breaking, so it can be played $1000 - 356 = 644$ more times.

DVD 2 has been played 135 times and can be played 1000 times before breaking, so it can be played $1000 - 135 = 865$ more times.

Now, we add the number of times each DVD can be played before breaking to the number of times each has already been played:

$$356 + 644 = 1000 \text{ (for DVD 1)}$$

$$135 + 865 = 1000 \text{ (for DVD 2)}$$

Now we add the two totals together:

$$1000 + 1000 = 2000$$

So, both DVDs can be played a total of **2000** times before they break. ✗

Role-Play Prompting: To find the total number of times both DVDs can be played before they break, we need to add the number of times each DVD has already been played and then subtract that total from the total number of times each DVD can be played.

The first DVD can be played 1000 times, and it has already been played 356 times. So, the number of times it can still be played is:

$$1000 - 356 = 644$$

The second DVD can be played 1000 times, and it has already been played 135 times. So, the number of times it can still be played is:

$$1000 - 135 = 865$$

Now, let's add the number of times each DVD can still be played:

$$644 + 865 = 1509$$

So, both DVDs can be played a total of **1509** times before they break. ✓

PERSONA SWITCH: To find the total number of times both DVDs can be played before they break, we need to add the number of times each DVD has already been played and then subtract that total from the total number of times each DVD can be played.

The first DVD can be played 1000 times, and it has already been played 356 times. So, the number of times it can still be played is:

$$1000 - 356 = 644$$

The second DVD can be played 1000 times, and it has already been played 135 times. So, the number of times it can still be played is:

$$1000 - 135 = 865$$

Now, we add the number of times each DVD can still be played:

$$644 + 865 = 1509$$

Therefore, both DVDs can be played a total of **1509** times before they break. ✓

Table 10: Full Example of GSM8K from Qualitative Analysis

Question: Take the last letters of each words in "Rich Parker Edward Milagros" and concatenate them.

Zero-Shot: To find the last letters of each word, we need to take the last letter of each word in the given phrase "Rich Parker Edward Milagros".

1. Rich - H
2. Parker - R
3. Edward - D
4. Milagros - S

Now, let's concatenate these letters: **H-R-D-S**. ✓

Role-Play Prompting: To find the last letters of each word, we'll look at the last letter of each name:

- Rich: H
- Parker: R
- Edward: D
- Milagros: S

Now, let's concatenate these letters: H-R-D-S. The resulting word is "**Hards**". ✗

PERSONA SWITCH: To find the last letters of each word, we need to take the last letter of each word in the given phrase "Rich Parker Edward Milagros"

1. Rich - H
2. Parker - R
3. Edward - D
4. Milagros - S

Now, let's concatenate these letters: **H-R-D-S**. ✓

Table 11: Full Example of Letter from Qualitative Analysis

Question: Alice, Bob, and Claire are holding a white elephant gift exchange. At the start of the event, they are each holding a present of a different color: Alice has a black ball, Bob has a green present, and Claire has an orange ball. As the event progresses, pairs of people swap gifts. First, Alice and Bob swap their gifts. Then, Bob and Claire swap their gifts. Finally, Claire and Alice swap their gifts. At the end of the event, Claire has the ... Which choice is true? Answer Choices: (A) black ball. (B) green present. (C) orange ball.

Zero-Shot: "Let's break down the gift swaps:

1. Alice and Bob swap gifts:
 - Alice has the green present (Bob's original gift)
 - Bob has the black ball (Alice's original gift)
2. Bob and Claire swap gifts:
 - Bob has the orange ball (Claire's original gift)
 - Claire has the green present (Bob's gift from the first swap)
3. Claire and Alice swap gifts:
 - Claire has the black ball (Alice's gift from the first swap)
 - Alice has the orange ball (Claire's gift from the second swap)

At the end of the event, Claire has the black ball.

The correct answer is (A) **black ball**. ✗

Role-Play Prompting: Let's follow the gift exchange process:

1. Alice and Bob swap gifts:
 - Alice has the green present (Bob's original gift)
 - Bob has the black ball (Alice's original gift)
 - Claire still has the orange ball
2. Bob and Claire swap gifts:
 - Bob has the orange ball (Claire's original gift)
 - Claire has the green present (Bob's gift after the first swap)
3. Claire and Alice swap gifts:
 - Claire has the black ball (Alice's gift after the first swap)
 - Alice has the orange ball (Claire's gift after the second swap)

At the end of the event, Claire has the black ball.

The correct answer is (A) **black ball**. ✗

PERSONA SWITCH: Let's follow the gift exchange process:

1. Alice and Bob swap gifts:
 - Alice has the green present (Bob's original gift)
 - Bob has the black ball (Alice's original gift)
 - Claire still has the orange ball (her original gift)
2. Bob and Claire swap gifts:
 - Alice still has the green present
 - Bob has the orange ball (Claire's original gift)
 - Claire has the black ball (Bob's gift from the first swap)
3. Claire and Alice swap gifts:
 - Alice has the black ball (Claire's gift from the second swap)
 - Claire has the green present (Alice's gift from the first swap)

So, at the end of the event, Claire has the green present.

The correct answer is (B) **green present**. ✓

Table 12: Full Example of Object from Qualitative Analysis