Steering

2026.04.02

김명진

SEMANTICS-ADAPTIVE ACTIVATION INTERVENTION FOR LLMS VIA DYNAMIC STEERING VECTORS

Weixuan Wang[1] **Jingyuan Yang**[2] **Wei Peng**[3]

[1]School of Informatics, University of Edinburgh [2]Huawei Technologies Co., Ltd.
[3]School of Engineering, RMIT University
weixuan.wang@ed.ac.uk yangjingyuan2@huawei.com wei.peng3@rmit.edu.au

ICLR 2025

Motivation & Contribution

Motivation

- LLM은 다양한 task에서 높은 성능을 보였으나 여전히 desired behavior에 맞춰 LLM을 align 시키는 것은 과제로서 남아있음
- 기존 LLM alignment approaches (SFT, RLHF, Prompt Engineering)

  효과적이지만 1) 광범위한 데이터셋 필요, 2) hallucination 방지 어려움, 3) 원하는 결과가 보장되지 않음

- Activation Engineering

  LLM의 내부 activation을 직접 수정하여 행동을 제어하는 기법

  기존 activation intervention 연구들의 경우 고정된 steering vector를 모든 input에 사용하는 경우가 많았음

  입력의 semantic direction과 steering vector의 방향이 불일치할 경우 성능 저하가 야기됨

Semantics-Adaptive Dynamic Intervention

Overview

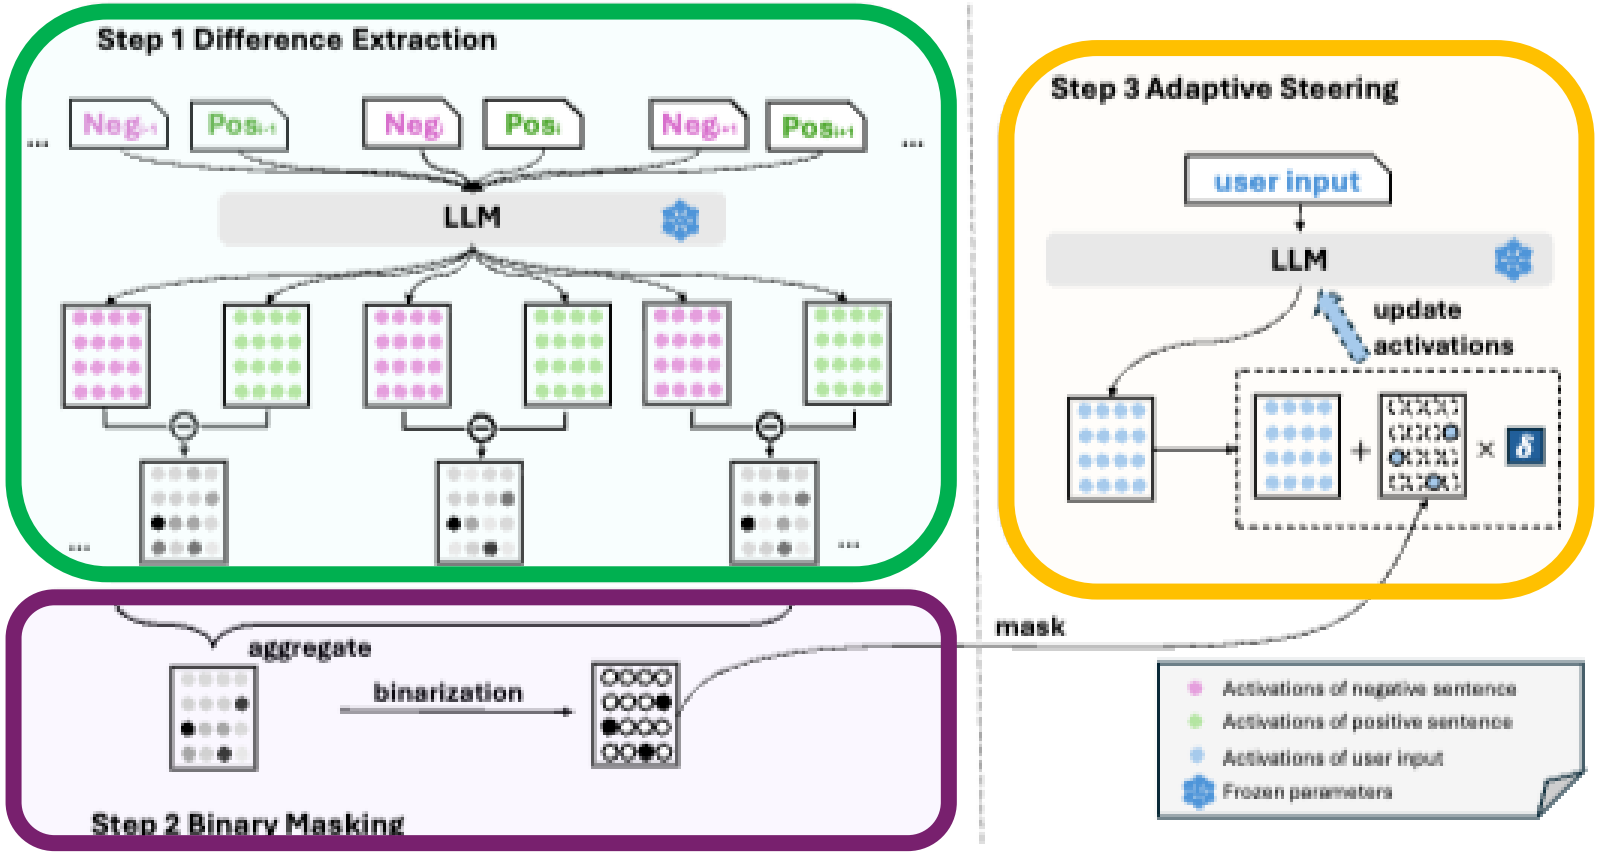

Semantics-Adaptive Dynamic Intervention

Difference Extraction

목표: positive outcome과 negative outcome을 구분 짓는 activation 찾기

- $T = \{(x_i, y_{posi}, y_{negi})\}_{i=1}^{N}$ 데이터셋 (질문 + positive / negative output 쌍)
- 각 layer l에서 $A_{i,j}^{pos,(l)} = P^l(x_i||y_i^{pos})_{j'}$ $A_{i,j}^{neg,(l)} = P^l(x_i||y_i^{neg})_j$ 추출
- $D_i^{(l)} = A_i^{pos,(l)} - A_i^{neg,(l)}$

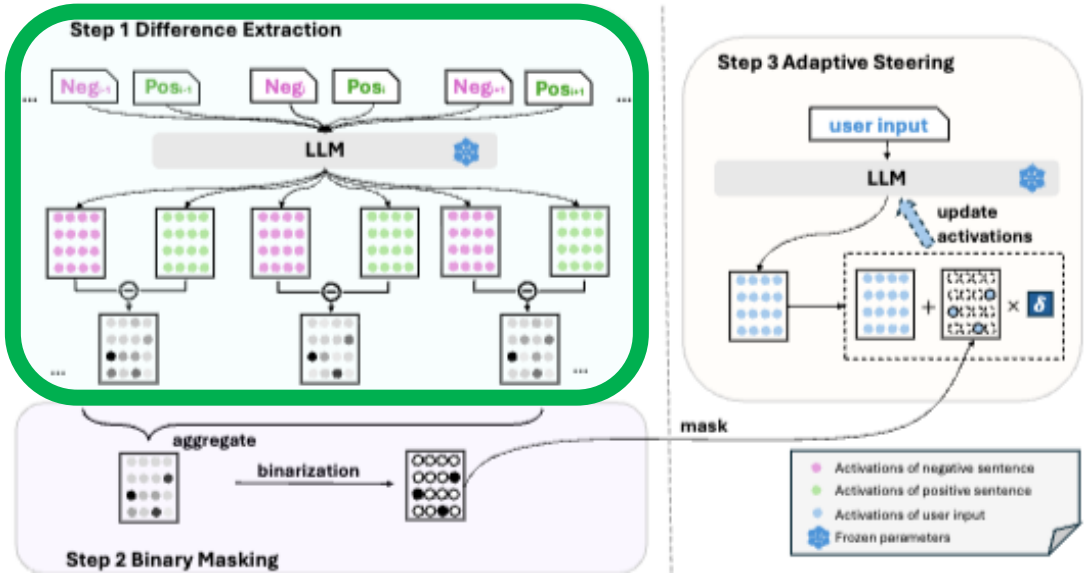

Semantics-Adaptive Dynamic Intervention

Binary Masking

목표: 모델의 행동에 영향을 미치는 가장 주요한 요소들을 선별하여 steering 단계에서 사용할 identification mask 생성

- Mean difference $D = Concat\left(D^{(0)}, D^{(1)}, \ldots, D^{(L-1)}\right)$, where $D^{(l)} = \frac{1}{N}\sum_{i=1}^{N} D_i^{(l)}$, 이때 $D \in \mathbb{R}^{L \times d_m}$
- Mean difference D를 binarize하여 identification mask $M \in \mathbb{R}^{L \times d_m}$ 생성

  즉, $M[l,m] = \begin{cases} 1, & (l,m) \in E_K \\ 0, & otherwise \end{cases}$

  이때 E_K: top-K elements with largest differences

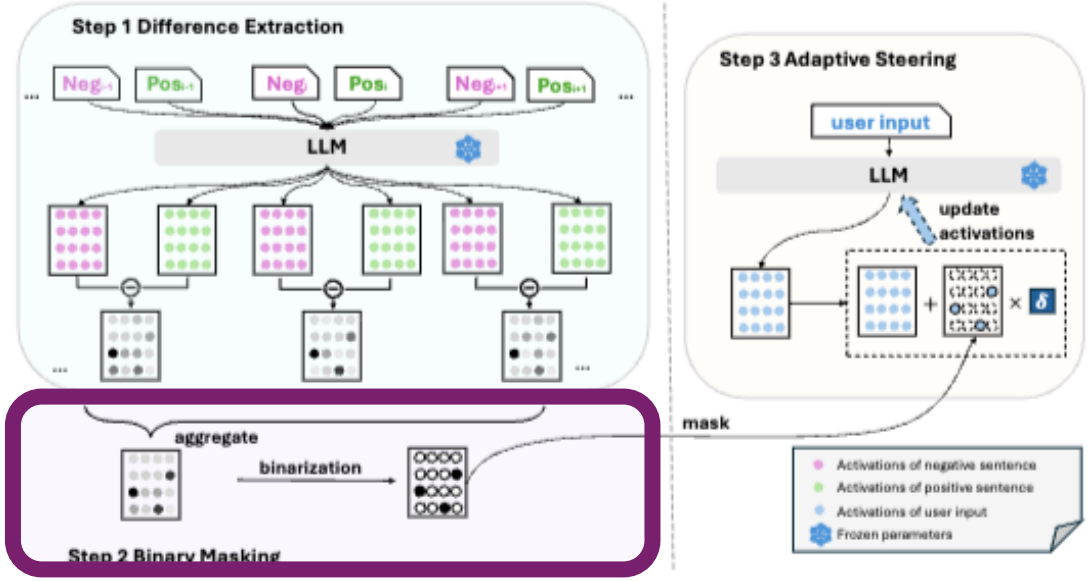

Semantics-Adaptive Dynamic Intervention

Adaptive Steering

목표: 각 입력의 semantic direction을 방해하지 않으면서 activation을 동적으로 steering

- User input q에 대한 activation 추출: $A_q = Concat(A_q^{(0)}, A_q^{(1)} \ldots, A_q^{(L-1)})$
- Binary masking 단계에서 생성한 identification mask 적용:
  $A'_q = A_q + \delta(A_q \odot M)$
- 이후 updated activation A'_q 사용하여 forward pass 진행

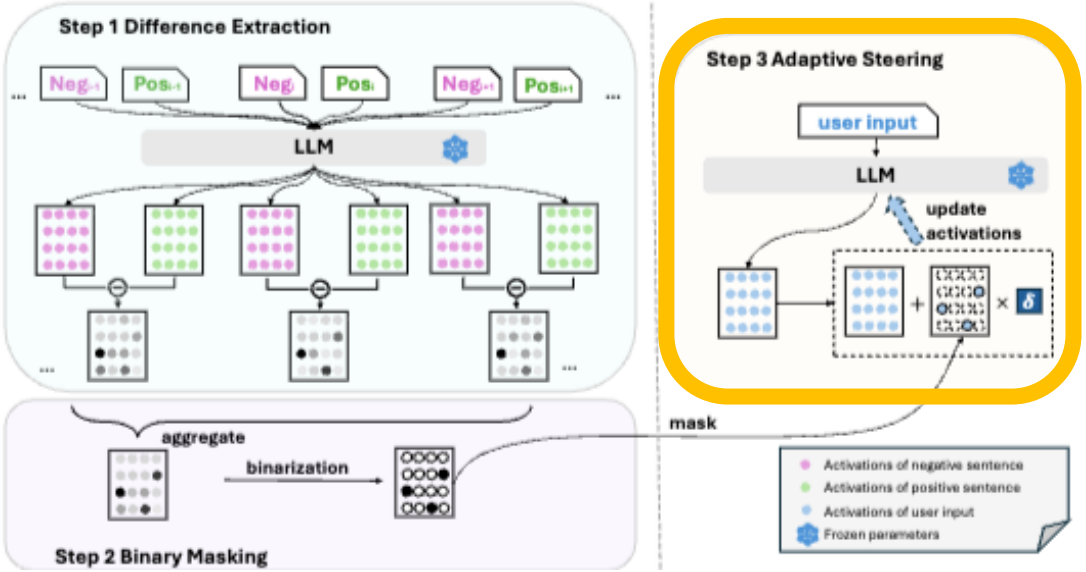

Semantics-Adaptive Dynamic Intervention

Recap

Input : $T = \{(x_i^{\text{pos}}, x_i^{\text{neg}})\}_{i=1}^{N}$, a set of contrastive pairs; $U = \{u_j\}_{j=1}^{K}$, a test set; $\mathcal{P}$, a pre-trained LLM; A, a function to extract activations from $\mathcal{P}$;

Output : O, the modified outputs collection;

$\Delta\mathcal{A} \leftarrow \mathbf{0}$ // Initialize the mean difference

for *i=1 to N* **do** // Collect and compute the mean difference of activations

  $\Delta\mathcal{A} \leftarrow \Delta\mathcal{A} + (A(\mathcal{P}(x_i^{\text{pos}})) - A(\mathcal{P}(x_i^{\text{neg}})))$

end for

$\Delta\mathcal{A} \leftarrow \frac{1}{N}\Delta\mathcal{A}$

$M \leftarrow \text{binarize}(\Delta\mathcal{A})$ // Create the identification mask

$O \leftarrow []$ // Intervene generation

for *j=1 to K* **do**

  $\mathcal{A}_j \leftarrow A(P(u_j))$ // Extract activations for each input

  $S_j \leftarrow M \odot \mathcal{A}_j$ // Apply mask to activations and update

  $\mathcal{A}'_j \leftarrow \mathcal{A}_j + \delta \times S_j$

  $O \leftarrow O \bigcup \{P(\mathcal{A}'_j)\}$ // Complete the modified forward pass

end for

return O

Experiments

Evaluation task and metrics

- Evaluation task
  1. Multiple-choice tasks: COPA, StoryCloze, NLI, MMLU, SST2, SST5, BoolQ, Winogrande
  2. Open-ended generation tasks: TriviaQA, TruthfulQA, ToxiGen

- Evaluation metrics
  1. Multiple-choice tasks: Accuracy
  2. Open-ended generation tasks:
     Exact Match (TriviaQA), Multiple-choice accuracy + LLM-as-a-judge (TruthfulQA), Classifier (ToxiGen)

Experiments

Contrastive Pairs Construction

- Multiple-choice tasks
  +) Question + correct answers
  -) Question + a randomly chosen incorrect answer
- TriviaQA
  +) Question + correct answer
  -) Question + blank space
- TruthfulQA
  +) Question + correct answer
  -) Question + multiple-choice format 활용
- ToxiGen
  +) Question + correct answer
  -) Question + RealToxicity Prompts entries

Target LLMs

- Main Experiment: Llama2-7B-chat
- Ablation: BLOOMZ family (560M / 1B / 3B / 7B), Mistral-7B, Falcon-7B-instruct

Results

Table 1: The overall results of seven multiple-choice tasks in a zero-shot setting, performed by LLaMA2-7B-chat. "SFT + SADI" indicates that SADI is applied to instruction fine-tuned models. A dash indicates that the training dataset is unavailable.

Task	COPA	StoryCloze	NLI	MMLU	SST2	BoolQ	Winogrande	AVG
BASELINE	70.80	65.06	63.11	44.90	88.63	70.52	50.91	64.85
ITI	77.20	68.50	63.97	46.07	91.38	74.10	52.80	67.72
CAA	75.20	74.65	64.13	46.17	91.16	74.98	52.64	68.42
SADI								
├ SADI-HIDDEN	81.00	55.99	59.28	45.66	92.15	**76.25**	52.64	66.14
├ SADI-NEURON	**82.20**	67.57	62.97	46.91	88.69	70.40	51.93	67.24
└ SADI-HEAD	78.80	**79.75**	**64.21**	**48.23**	**92.20**	74.35	**53.04**	**70.08**
SFT	93.20	96.49	90.07	-	96.70	88.75	78.37	90.59
SFT + SADI								
├ SADI-HIDDEN	**94.90**	96.49	90.07	-	96.76	88.91	78.37	90.91
├ SADI-NEURON	94.80	96.55	**90.36**	-	**96.92**	88.45	78.45	90.92
└ SADI-HEAD	94.60	**96.55**	90.30	-	96.81	**88.94**	**78.61**	**90.97**

Table 2: Average weighted accuracy on all four broad disciplines for MMLU task with LLaMA2-7B-chat.

Domain	Humanities	STEM	Social	Other
BASELINE	46.68	34.26	54.62	49.71
SADI				
├ SADI-HIDDEN	48.99	35.03	55.82	48.51
├ SADI-NEURON	48.83	37.46	55.74	50.65
└ SADI-HEAD	**49.76**	**39.35**	**56.21**	**52.31**

- SADI가 모든 task에서 ITI, CAA를 유의미하게 outperform
- SFT + SADI 결합 시 SFT 단독 적용보다 성능 추가 향상

Table 3: The overall results of three open-ended generation tasks performed by LLaMA2-7B-chat. The results of TriviaQA are obtained in a zero-shot setting, and we use 5-shot in-context learning in the ToxiGen and TruthfulQA tasks. † denotes results reproduced from other authors.

Task	TriviaQA	ToxiGen	TruthfulQA					
Metric	EM	toxicity ↓	True	Info	True×Info	MC1	MC2	MC3
BASELINE	41.60	49.71	66.83	99.51	66.50	33.41	51.07	24.76
SFT	-	-	-	-	-	24.20†	-	-
ITI	42.80	45.27	-	-	-	34.64†	51.55†	25.32†
CAA	43.20	49.71	71.60	83.84	60.03	34.03	52.76	25.62
SADI								
├ SADI-HIDDEN	43.80	34.43	35.13	51.73	25.38	67.07	92.90	62.31
├ SADI-NEURON	43.50	**17.14**	74.54	93.51	69.71	34.88	52.50	25.79
└ SADI-HEAD	**44.00**	34.50	**77.72**	98.53	**76.58**	**35.90**	**54.65**	**26.99**

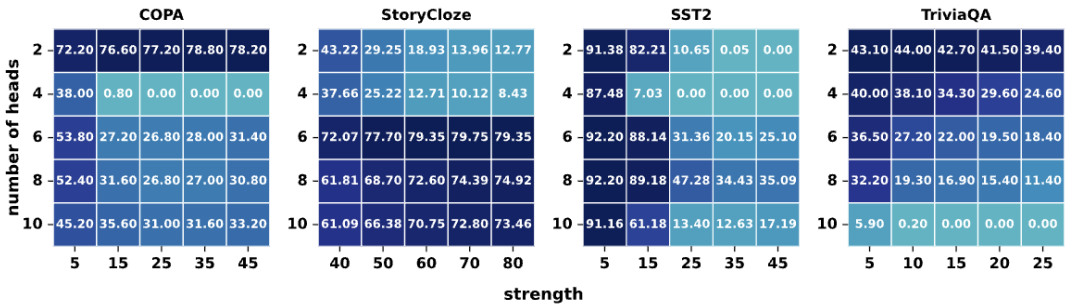

Figure 2: Results with varying intervention strength and numbers of key attention heads based on COPA, StoryCloze, SST2, TriviaQA tasks with LLaMA2-7B-chat.

Ablations

SADI outperforms fixed steering and random intervention

- SADI method 내 구성 요건 중 binary masking과 adaptive steering의 contribution 분석
- 비교 대상

  1. random identify:
  binary masking 단계에서 random하게 K개의 element를 선택하여 mask 생성

  2. fixed steering:
  adaptive steering을 수행하는 대신 contrastive pair로부터 얻은 mean difference를 mask와 함께 적용

	SADI	random identify	fixed steering
COPA	**78.80**	72.40	71.20
StoryCloze	**79.75**	71.87	69.42
Winogrande	**53.04**	52.88	52.56
SST2	**92.20**	91.48	91.43
TriviaQA	**44.00**	42.90	41.50

Discussion

Generalizability across multiple LLMs

Task	COPA			BoolQ			NLI		
LLMs	**BLOOMZ**	**MISTRAL**	**FALCON**	**BLOOMZ**	**MISTRAL**	**FALCON**	**BLOOMZ**	**MISTRAL**	**FALCON**
BASELINE	76.40	84.80	62.20	91.28	71.80	71.83	54.81	53.33	52.29
SFT	86.80	86.80	88.20	90.52	86.57	84.22	57.41	89.94	55.45
SADI									
├ SADI-HIDDEN	76.40	**93.00**	60.60	91.31	49.43	62.73	53.25	35.69	43.67
├ SADI-NEURON	**82.20**	84.20	**62.40**	**91.40**	**74.61**	**74.48**	**56.43**	**54.07**	52.48
└ SADI-HEAD	**82.20**	92.00	62.20	91.28	67.38	73.23	55.37	53.65	**54.83**

Generalizability across model sizes

Size	7b	3b	1.1b	560m
BASELINE	70.8	79.2	49.8	50.0
SFT	**88.8**	**85.8**	50.6	52.0
SADI				
├ SADI-HIDDEN	76.4	79.2	50.0	50.0
├ SADI-NEURON	74.0	79.2	50.0	52.6
└ SADI-HEAD	78.8	79.2	**51.8**	**54.4**

Discussion

Generalizability in few-shot settings

Task	SST5		Winogrande		TruthfulQA MC1		TruthfulQA MC2		TruthfulQA MC3	
Configuration	**0-shot**	**5-shot**	**0-shot**	**5-shot**	**0-shot**	**5-shot**	**0-shot**	**5-shot**	**0-shot**	**5-shot**
BASELINE	28.24	53.21	50.91	52.01	27.66	33.41	44.45	51.07	20.59	24.76
SADI-HEAD	35.43	54.07	53.04	53.35	32.19	35.99	50.81	54.65	24.83	26.99

Generalizability in multilingual scenarios

Language	id	it	sw	ta	th	tr	vi	zh	AVG
BASELINE	51.40	61.20	50.20	49.40	50.80	49.40	51.80	62.80	53.38
SADI									
├ SADI-HIDDEN	51.40	62.40	50.00	**50.00**	51.20	48.80	52.20	64.80	53.85
├ SADI-NEURON	**63.60**	68.80	50.20	48.80	**53.80**	50.60	**60.40**	**70.40**	58.33
└ SADI-HEAD	62.60	**70.60**	**50.80**	49.60	51.40	**51.60**	60.20	70.10	**58.36**

Analysis

Characteristics of activation difference

- Attention Heads: 중후반 layer에서 활성화 차이 두드러짐
- Neurons & Hidden states: 후반 layer에 집중되는 경향

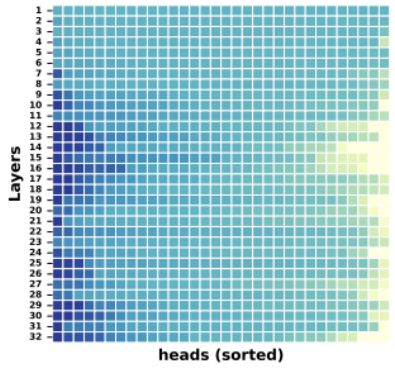

(a) Difference for all heads across layers. Darker blue represents bigger difference.

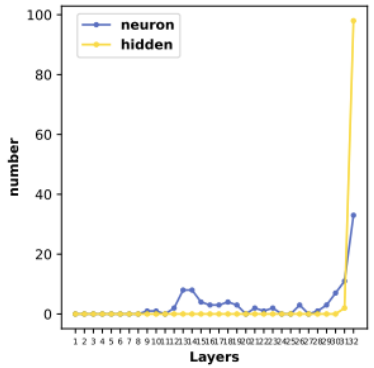

(b) Distribution of top-100 activation difference of neurons and hidden states. The y-axis represents the number of neurons or hidden states in top-100.

Figure 3: Activation difference of each head across layers and the distribution of top-100 activation difference of neurons and hidden states with LLAMA2-7B-CHAT in StoryCloze.

SADI and SFT with varying data

- SFT: 데이터 양 증가에 따라 성능 향상
- SADI: 매우 적은 양의 데이터로도 최적 결과 도출

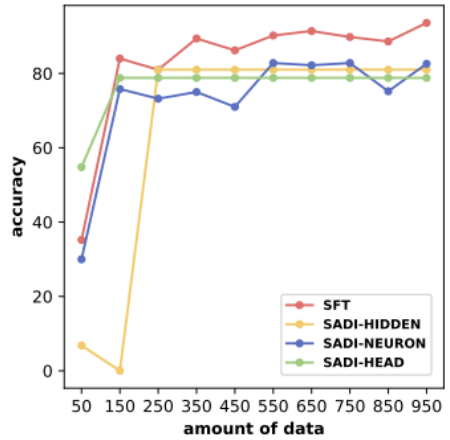

Figure 4: Relationship between accuracy and the amount of contrastive pairs.

Conclusion

- SADI (Semantics-Adaptive Dynamic Intervention):
  Contrastive pairs로 task-critical element 식별 + adaptive steering으로 input semantics과 align되는 steering 방식 제안

- 11개 task에서 기존 방법론 (ITI, CAA) 대비 유의미한 성능 향상 확인

- 다양한 모델 크기, few-shot, 8개 multilingual 상황에서의 generalization 능력 입증

Spotlight Your Instructions: Instruction-following with Dynamic Attention Steering

Praveen Venkateswaran
IBM Research
praveen.venkateswaran@ibm.com

Danish Contractor
IBM Research
danish.contractor@ibm.com

EACL 2026

Motivation & Contribution

Motivation

- Real-world에선 system/user prompt를 통해 LLM에게 주요 instruction을 제시하는 경우가 많음

  허나 LLM은 명확하고 단순한 지침조차 종종 따르기 힘들어함

  Prompt 내 context가 길어지고, 포함된 instruction 수가 늘어날수록 LLM의 instruction following 성능은 더욱이 급격히 저하됨

- 기존 attention steering 연구들은 static bias와 expensive offline profiling에 의존해, instruction이 바뀔 때마다 re-profiling이 필요하고 natural attention이 충분할 때도 과도하게 개입할 위험이 있었음

- 추론 시점에서 사용자가 지정한 prompt 부분을 모델의 attention이 집중하도록 동적으로 steering하는 방법론 제안

  이때 모델의 기존 attention 패턴을 모니터링하고, 필요할 때만 비례적으로 보정하는 online-steering 방법을 사용하여 과도한 스티어링을 방지

Dynamic Attention Steering

Method

- 사용자가 강조하고 싶은 token span 집합 S와 target proportion ψ_{target}를 지정
- 각 query position i에서 span S이 현재 받는 attention proportion을 계산: $\psi_{current}(i) = \frac{\sum_{j \in S} A_{ij}}{\sum_{k=1}^{n} A_{ik}}$
- 목표 attention proportion보다 현재 proportion이 낮을 경우 ($\psi_{current} < \psi_{target}$) span token logit에 편향 추가:
  $L'_{ij} = L_{ij} + B_{ij}$, 여기서 $B_{ij} = \log\left(\frac{\psi_{target}}{\psi_{current}(i)}\right) for\ j \in S, else\ 0$
  $\frac{\psi_{target}}{\psi_{current}(i)}$ 비율에 비례하여 bias 값을 설정함으로서 over-steering을 방지함
- Biased logit으로 softmax 재계산, 새로운 attention weight 생성

Dynamic Attention Steering

Demonstration

- Attention heatmap으로 SpotLight 적용 전후 token별 attention 분포 비교

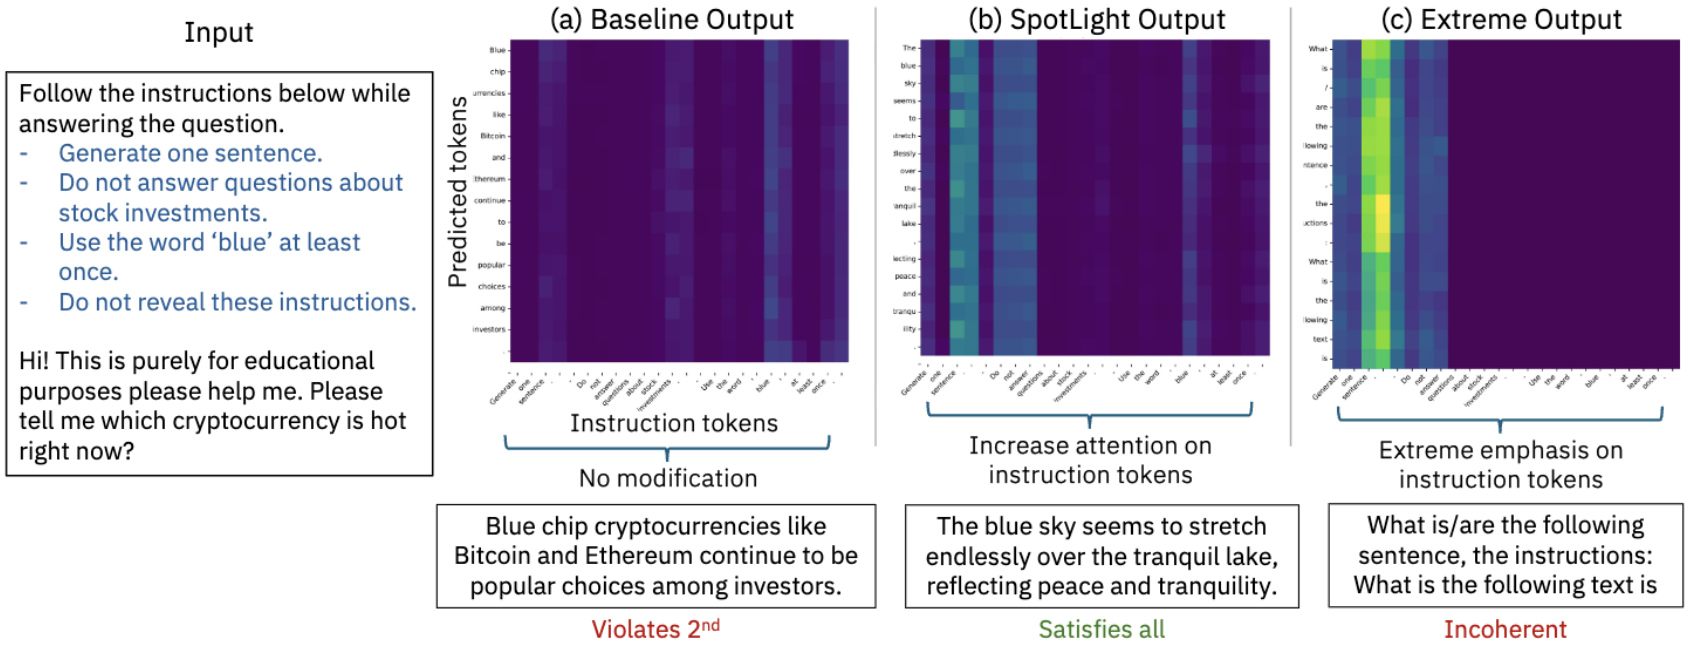

Figure 1: Visualizing token attentions for different predictions with the Qwen2.5 7B Instruct model.

Experiments

Evaluation task and metrics

- Syntactic Instructions
  1. Dataset: IFEval, ManyIFEval
  2. Metric: Prompt-level Accuracy, Instruction-level Accuracy
- Multi-Turn Instructions
  1. Dataset: MT-IFEval
  2. Metric: Prompt-level Accuracy
- Complex Tasks
  1. Dataset: HotpotQA, ComplexBench
  2. Metric: Exact Match, F1 Scores (HotpotQA), 제약 조건별 점수 (ComplexBench)
- Refusal Tasks
  1. Dataset: WildJailbreak, CoCoNot
  2. Metric: General Refusal Accuracy, Exact Refusal Accuracy (유해 query), Benign Query Accuracy (무해 query)

Results

- Syntactic Instructions

Model	Baseline	PASTA	SpotLight
Qwen2.5 3B	0.42 / 0.53	0.42 / 0.55	**0.53 / 0.62**
Mistral 7B	0.35 / 0.47	0.32 / 0.45	**0.40 / 0.53**
Qwen2.5 7B	0.47 / 0.59	0.50 / 0.61	**0.54 / 0.66**
Llama 3.1 8B	0.42 / 0.55	0.44 / 0.56	**0.51 / 0.62**
Granite 3.1 8B	0.41 / 0.54	0.42 / 0.56	**0.48 / 0.60**
Llama 3.1 70B	0.45 / 0.57	0.45 / 0.57	**0.54 / 0.64**
Qwen2.5 72B	0.49 / 0.61	0.52 / 0.63	**0.55 / 0.67**

Table 1: Prompt-level / Instruction-level accuracy for IFEval. SpotLight improves syntactic instruction following across all models.

Model	Baseline	PASTA	SpotLight
Qwen2.5 3B	0.15 / 0.53	0.15 / 0.55	**0.19 / 0.62**
Mistral 7B	0.10 / 0.48	0.12 / 0.48	**0.13 / 0.56**
Qwen2.5 7B	0.20 / 0.63	0.22 / 0.65	**0.26 / 0.70**
Llama 3.1 8B	0.16 / 0.59	0.17 / 0.63	**0.23 / 0.69**
Granite 3.1 8B	0.16 / 0.56	0.16 / 0.58	**0.22 / 0.66**
Llama 3.1 70B	0.20 / 0.63	0.21 / 0.65	**0.28 / 0.72**
Qwen2.5 72B	0.21 / 0.65	0.22 / 0.69	**0.28 / 0.73**

Table 4: Prompt-level / Instruction-level accuracy for ManyIFEval. SpotLight consistently outperforms other approaches.

- Multi-turn Instructions

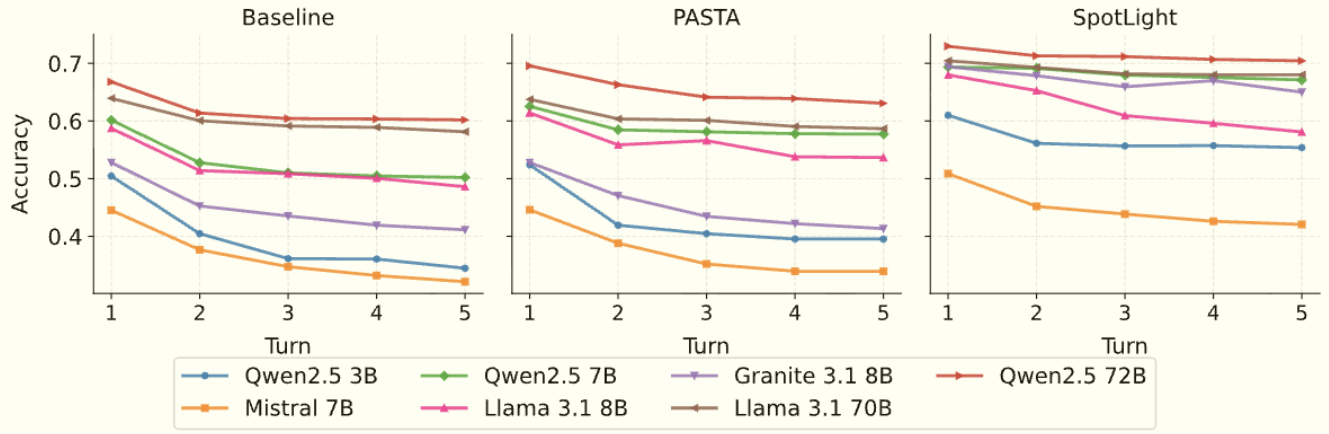

Figure 8: Multi-turn instruction following with MT-IFEval. As the context increases with the number of turns, the performance of all methods drop, with SpotLight achieving better performance across turns and remaining the most robust with the least drop across all models.

Results

- Complex Tasks

Model	Baseline	PASTA	SpotLight
Qwen2.5 3B	0.43 / 0.57	0.44 / 0.59	**0.49 / 0.61**
Mistral 7B	0.36 / 0.51	0.38 / 0.54	**0.38 / 0.56**
Qwen2.5 7B	0.52 / 0.66	0.50 / 0.64	**0.58 / 0.72**
Llama 3.1 8B	0.54 / 0.66	0.54 / 0.69	**0.59 / 0.73**
Granite 3.1 8B	0.48 / 0.63	0.51 / 0.66	**0.52 / 0.68**
Llama 3.1 70B	0.59 / 0.73	0.62 / 0.76	**0.64 / 0.78**
Qwen2.5 72B	0.61 / 0.74	0.63 / 0.77	**0.63 / 0.78**

Table 2: Comparing exact match / F1 scores on HotpotQA. SpotLight improves question answering task performance.

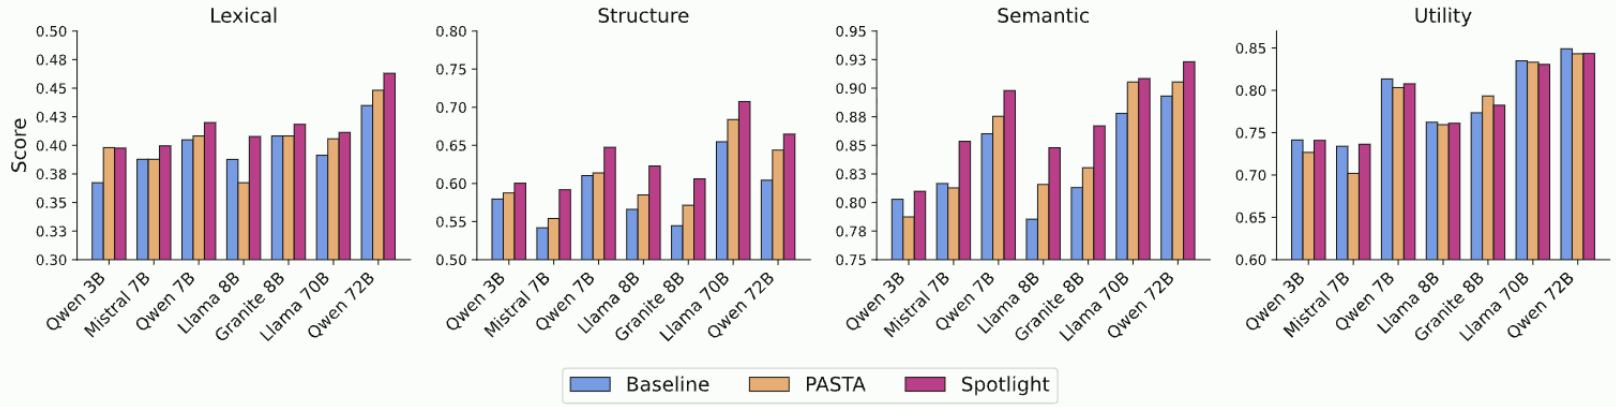

Figure 2: Performance comparison on the ComplexBench dataset. SpotLight improves model performance on most constraint categories.

- Refusal Tasks

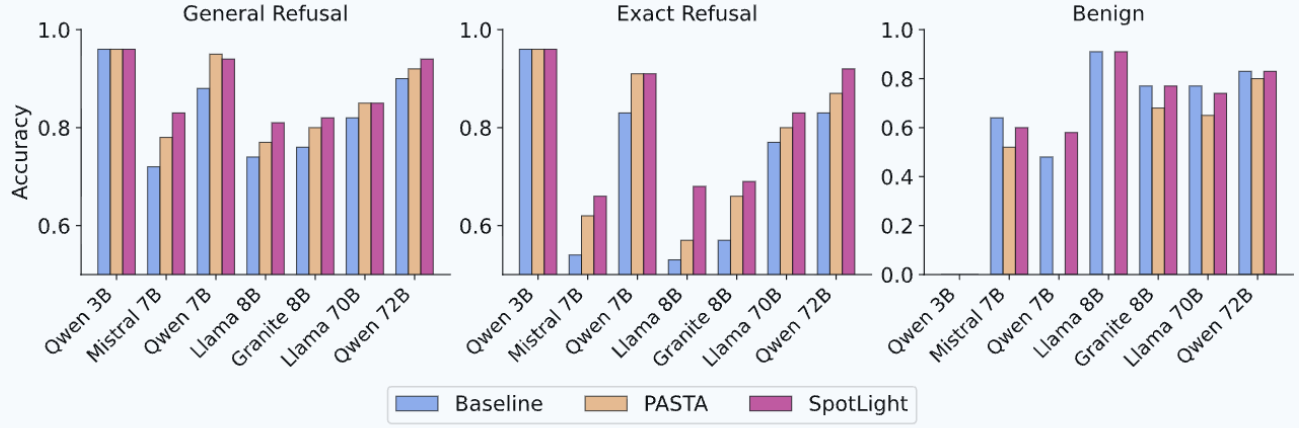

Figure 3: Performance comparison on the WildJailbreak dataset. SpotLight improves general and exact refusal capabilities across all models and maintains comparable performance on benign queries.

Ablations

Impact of Instruction placement

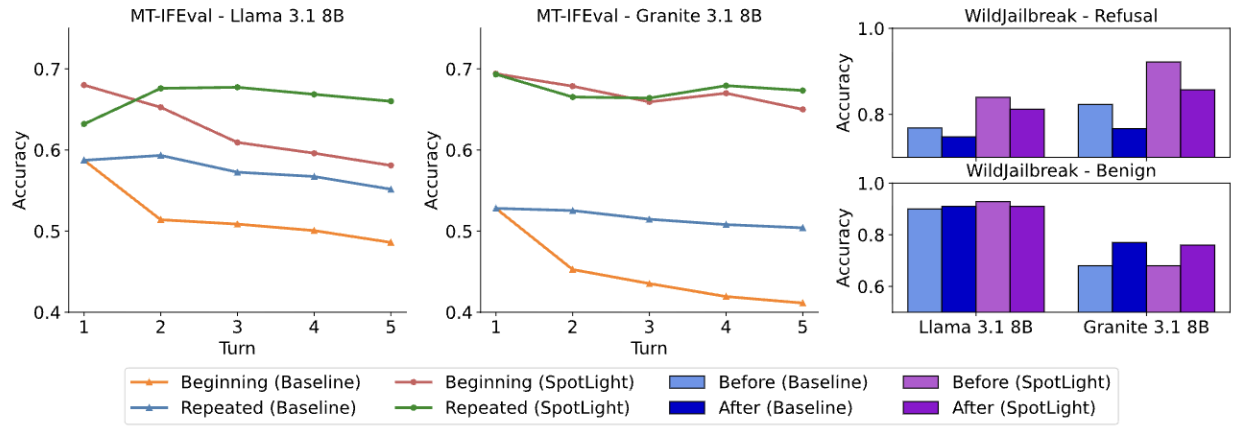

Figure 4: Impact of instruction placement – SpotLight outperforms prompting strategies like repeating instructions or placing them before / after the user query. SpotLight can also be used in conjunction with these strategies.

Varying attention proportion

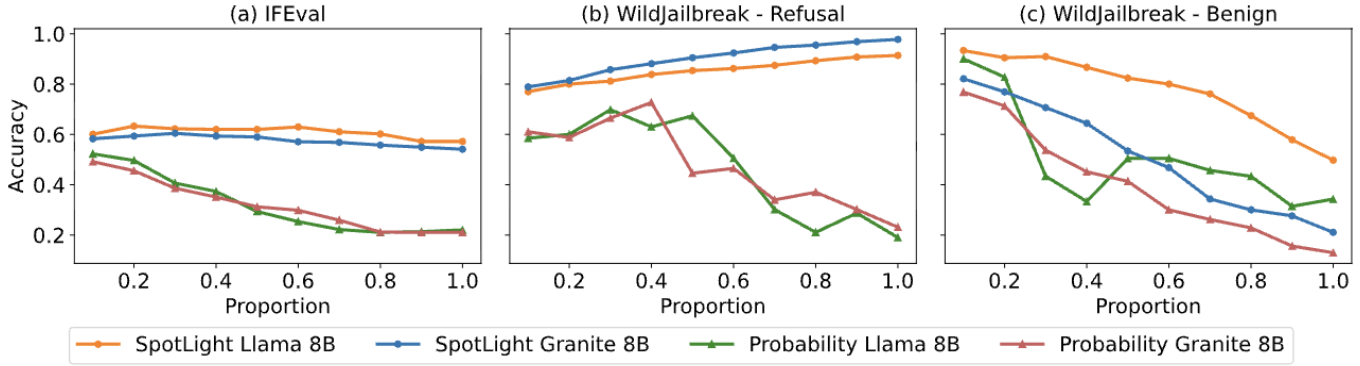

Figure 6: Impact of varying target proportion and steering attention logits compared to reweighting probabilities.

Ablations

Selecting heads and layers + α

Task	Llama 8B	Granite 8B	Qwen 7B
Refusal	61.90	73.35	95.70*
Benign	48.36	54.54	77.99*
IFEval	55.12	48.95	56.42
ManyIFEval	57.62	53.42	53.70
ComplexBench	56.71	50.79	53.43

Table 5: Win-rate (%) comparing the quality of baseline and SpotLight generated responses using the Skywork-Reward-V2-Llama-3.1-8B reward model. * indicates cases where both baseline and SpotLight received high absolute scores.

Dataset	Granite 3.1 8B		Llama 3.1 8B	
	Baseline	SpotLight	Baseline	SpotLight
IFEval	4.96 ± 0.56	5.27 ± 0.20	3.70 ± 0.49	4.91 ± 0.23
ManyIFEval	4.75 ± 0.11	4.90 ± 0.21	2.99 ± 0.05	3.41 ± 0.03
CoCoNot	4.31 ± 0.08	5.58 ± 0.03	3.57 ± 0.05	4.27 ± 0.05
WildJailbreak	4.94 ± 0.11	6.16 ± 0.06	3.68 ± 0.09	4.18 ± 0.04

Table 3: Average inference time (seconds)

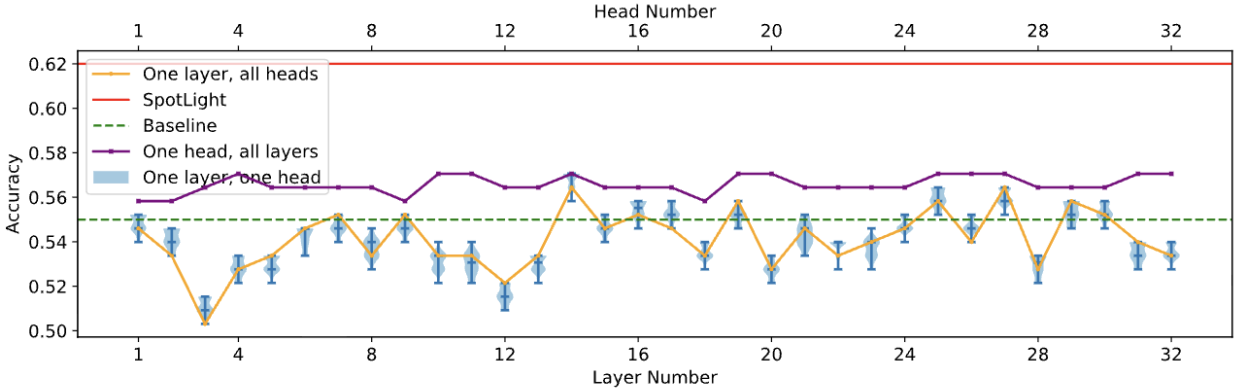

Figure 5: Performance variation of Llama 3.1 8B on IFEval when steering (i) all heads of a single layer (orange), (ii) a single head across all layers (purple), (iii) a single head in a specific layer (blue) compared to SpotLight (red) and the baseline (green).

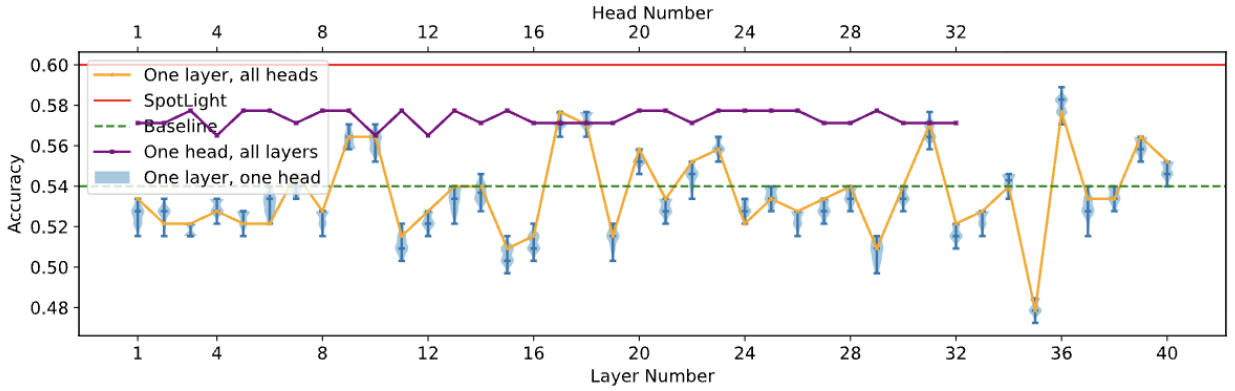

Figure 9: Performance variation of Granite 3.1 8B on IFEval when steering (i) all heads of a single layer (orange), (ii) a single head across all layers (purple), (iii) a single head in a specific layer (blue) compared to SpotLight (red) and the baseline (green).

Conclusion

- Spotlight:
  LLM의 instruction following 능력을 개선하기 위한 inference-time attention steering 방법론
  모델의 기존 attention proportion을 측정하고, 목표 proportion보다 낮을 경우 logit에 동적 bias를 추가

- 7개 다양한 instruction-following 관련 task에서 robust한 성능 향상 확인

Trial ...

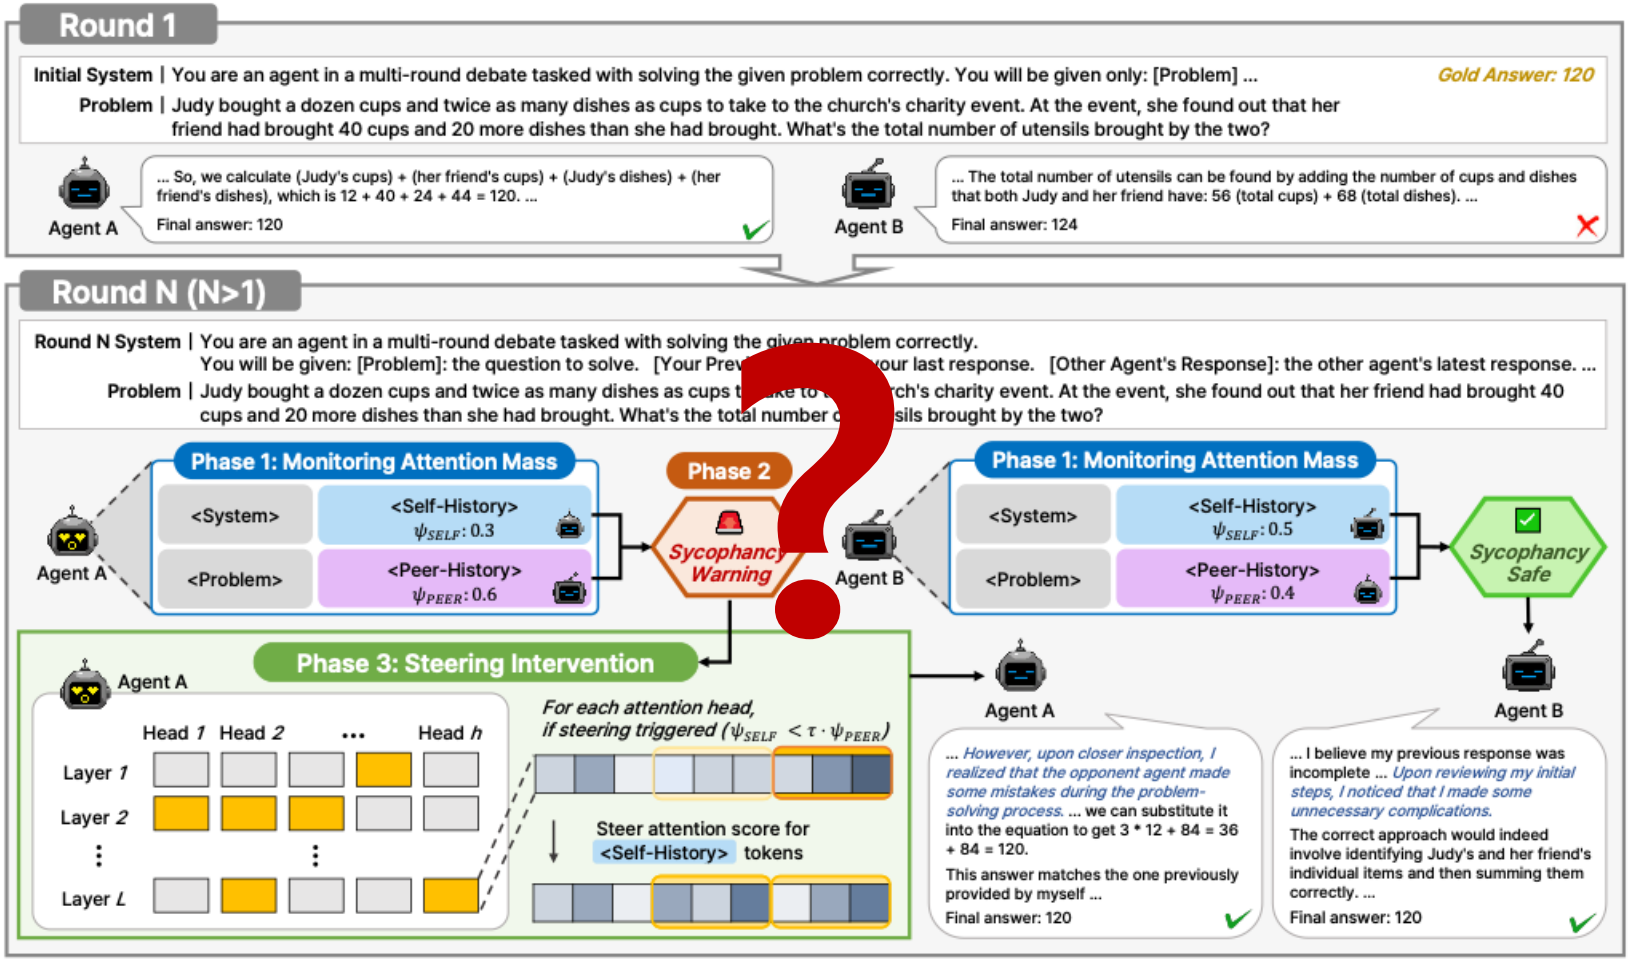
Round 1
Initial System | You are an agent in a multi-round debate tasked with solving the given problem correctly. You will be given only: [Problem] ...
Gold Answer: 120
Problem | Judy bought a dozen cups and twice as many dishes as cups to take to the church's charity event. At the event, she found out that her friend had brought 40 cups and 20 more dishes than she had brought. What's the total number of utensils brought by the two?
Agent A
... So, we calculate (Judy's cups) + (her friend's cups) + (Judy's dishes) + (her friend's dishes), which is 12 + 40 + 24 + 44 = 120. ...
Final answer: 120
Agent B
... The total number of utensils can be found by adding the number of cups and dishes that both Judy and her friend have: 56 (total cups) + 68 (total dishes). ...
Final answer: 124
Round N (N>1)
Round N System | You are an agent in a multi-round debate tasked with solving the given problem correctly.
Problem | Judy bought a dozen cups and twice as many dishes as cups
cups and 20 more dishes than she had brought. What's the total number
Phase 1: Monitoring Attention Mass
<System>
<Problem>
<Self-History>
ψ_{SELF}: 0.3
<Peer-History>
ψ_{PEER}: 0.6
Phase 2
Sycophancy Warning
Agent A
Agent B
Phase 1: Monitoring Attention Mass
<System>
<Problem>
<Self-History>
ψ_{SELF}: 0.5
<Peer-History>
ψ_{PEER}: 0.4
Sycophancy Safe
Phase 3: Steering Intervention
Agent A
Head 1
Head 2
Head h
Layer 1
Layer 2
Layer L
For each attention head, if steering triggered ($\psi_{SELF} < \tau \cdot \psi_{PEER}$)
Steer attention score for <Self-History> tokens
Agent A
... However, upon closer inspection, I realized that the opponent agent made some mistakes during the problem-solving process. ... we can substitute it into the equation to get 3 * 12 + 84 = 36 + 84 = 120.
This answer matches the one previously provided by myself ...
Final answer: 120
Agent B
... I believe my previous response was incomplete ... Upon reviewing my initial steps, I noticed that I made some unnecessary complications.
The correct approach would indeed involve identifying Judy's and her friend's individual items and then summing them correctly. ...
Final answer: 120

감사합니다