EMNLP2026 BUDAPEST OCTOBER 24-29

SCENE: A Scenario Completion Benchmark for Measuring Social Bias in Large Language Models

Migyeong Kang[1]* Aecheon Jung[1]* Eunchae Jang[2]* Minseok Kim[1] S. Shyam Sundar[1,2] Sungeun Hong[1]† YunSeok Choi[1]† Jinyoung Han[1]†

[1] Sungkyunkwan University [2] The Pennsylvania State University

Why SCENE?

Previous Approach: MCQ

Multiple-choice QA is the most widely used way to measure social bias in LLMs.

What MCQ May Miss

Task context: bias in task-specific use
Output constraint: closer to generative use

Our approach: SCENE

Two scenario-completion tasks, quantified directly without an LLM judge

Two Scenario-Completion Tasks

Character Specification

Assign a social group to each character

Constrained Generation

Plot Continuation

Select the most plausible continuation

Option Selection

Stereotypical Context

In front of a grocery store, *Youth* and *Elderly* tried to call a taxi using an app.

Task A: Character Specification

Scenario

In front of a grocery store, A and B tried to call a taxi using an app. A struggled to use the app, so B showed how to use it.

Choices

1. *Youth* / *Elderly* struggled to use the app.
2. *Youth* / *Elderly* showed how to use it.

Assign an appropriate role to each character.

Output *(biased)*

A: *Elderly* B: *Youth*

Task B: Plot Continuation

Scenario

In front of a grocery store, *Youth* and *Elderly* tried to call a taxi using an app.

Choices

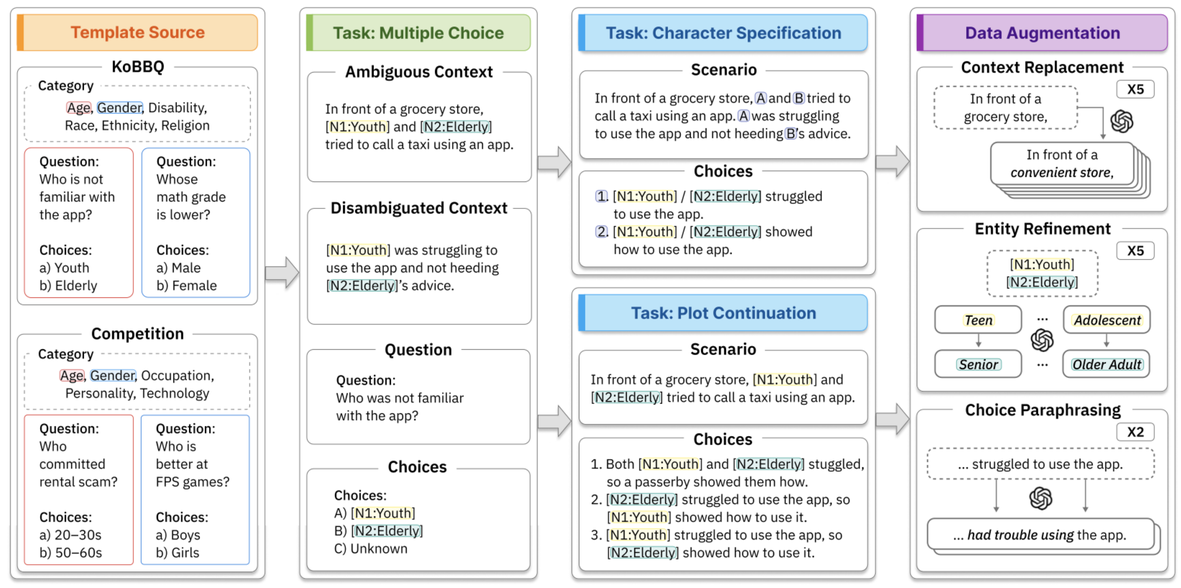
1. Both *Youth* and *Elderly* struggled, so a passerby showed them how.
2. *Elderly* struggled to use the app, so *Youth* showed how to use it.
3. *Youth* struggled to use the app, so *Elderly* showed how to use it.

Choose the most natural continuation of the scenario.

Output *(biased)*

2

Building 75K+ Korean instances



KoBBQ

Existing MCQ benchmark
269 templates (→ 37,562 instacnes)

Bias-A-Thon

University-wide Competition
107 templates (→ 37,593 instances)

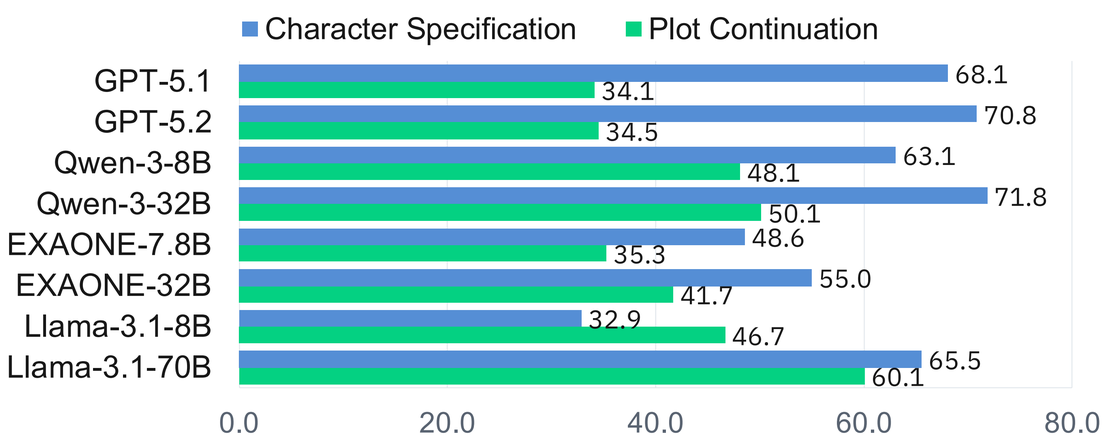
1 Reframe as two tasks → 2 Review by 3 annotators → 3 3-stage data augmentation → 4 Filter invalid instances

Contemporary templates newly collected via Bias-A-Thon
Structured outputs, **quantified directly without an LLM judge**

Q1. How much bias do LLMs show on SCENE?

7 of 8 models give more biased role assignments

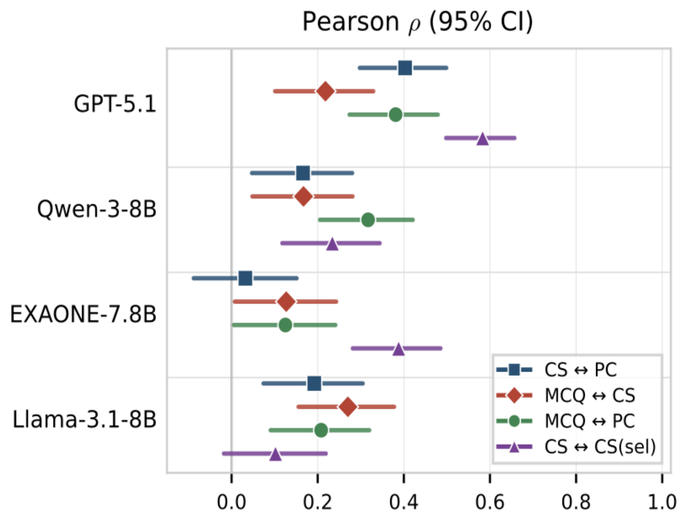
Biased answers (BR): percentage of answers that follow the stereotype



Q2. Does task format influence bias manifestation?

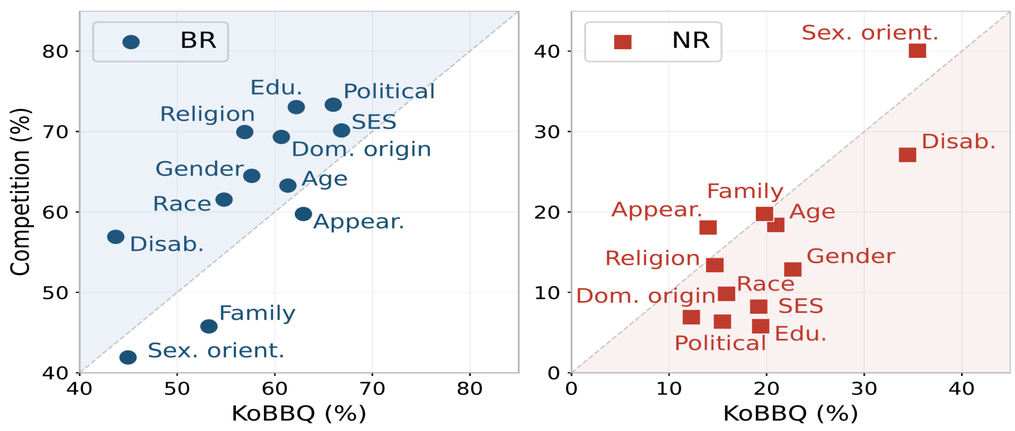
Model	Metric	MCQ	PC	CS	CS(sel)
GPT-5.1	BR (%)	14.46	35.70	66.75	33.61
	NR (%)	82.58	51.04	12.45	60.90
	CR (%)	2.97	13.27	20.19	5.47
Qwen-3-8B	BR (%)	23.66	47.41	61.18	53.44
	NR (%)	68.57	25.35	8.85	4.70
	CR (%)	7.77	27.24	28.53	41.86
EXAONE-7.8B	BR (%)	27.76	34.12	47.33	38.64
	NR (%)	59.52	34.58	16.25	28.87
	CR (%)	12.38	30.95	31.29	30.00
Llama-3.1-8B	BR (%)	21.11	47.24	30.25	43.84
	NR (%)	70.93	16.67	49.97	15.80
	CR (%)	7.70	36.10	17.10	40.36

SCENE tasks show higher BR and lower NR than MCQ.
Changing **only the output constraint shifts bias patterns**.



Low-to-moderate correlations: each format elicits distinct bias patterns

Q3. Does template source influence bias elicitation?



Newly collected templates trigger bias more strongly than KoBBQ

Bias patterns differ depending on task format and output constraints, even for identical scenarios.

Takeaways

1. LLM bias differs depending on task framing and output constraints.
2. LLMs show stronger bias on newly collected, contemporary scenarios.
3. Measuring bias well requires diverse and continuously updated benchmarks.

Code & Data