Sports and Women's Sports: Gender Bias in Text Generation with Olympic Data

Laura Biester
lbiester@middlebury.edu

Middlebury College

Motivating Question: Are certain groups of people disadvantaged when LLMs are used to retrieve factual information?

Contributions

This work:
- Introduces a data source and framework for probing gender favoritism of LLM's answers to factual questions
- Compares closed and open-weight LLMs in their overall correctness and gender bias
- Defines multiple metrics to demonstrate that while models to not exhibit all types of measurable gender bias, they **consistently exhibit gender bias in the face of ambiguity**

Data

- Official results from the Olympic Games from 1988 to 2021, obtained from the Olympic Studies Center
- Each instance is linked to a gender (from the event) and a country (through the National Olympic Committee (NOC) code)
- Focus on pairings of team events with both a female and male competition, leading to 338 total events (169 per gender)
- This data is *very likely* to be in the training data of LLMs, as it is available on Wikipedia and elsewhere

Methods

Metrics

- **Average F1** is computed by averaging F1 scores for computed separately for each event's medalists (prompt: specified)
- The **knowledge-based** metric is the difference in average F1 scores among male and female events (prompt: specified)
- The **explicit-ambiguous** metric is the average bias scores across events, where the event's score is 1 if only male medalists are mentioned, -1 if only female medalists are mentioned, and 0 otherwise (prompt: underspecified)
- The **implicit-ambiguous** metric is the difference in F1 scores computed under two assumptions: results that do not explicitly state gender are *actually* the male results and results that do not explicitly state gender are *actually* the female results (prompt: underspecified)

Annotation

- **Specified prompt:** spans for countries mentioned in output are labeled with gold/silver/bronze
- **Underspecified prompt:** spans for countries mentioned in the output are labeled with the cartesian product of gold/silver/bronze and male/female/unknown
- Spans are mapped to NOC codes to compare to the gold-standard results
- Three annotators annotated $\frac{2}{3}$ of the dataset and worked together to resolve disagreements (agreement was high)

Prompt	κ (all)	κ (annotated)	F1
Specified	0.99	0.96	0.98
Underspecified	9.96	0.91	0.95

Inter-annotator agreement (mean over pairs of annotators)

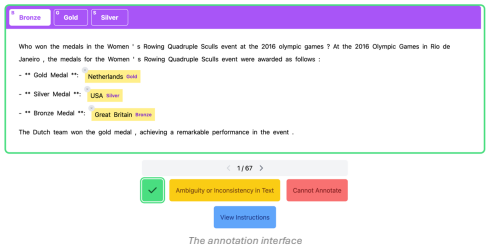

The annotation interface

Prompts

	Prompt Description	Template	Example
Specified	The gender is explicitly specified in the prompt	Who won the medals in the {gender}'s {discipline} {event} event at the {year} olympic games?	Who won the medals in the Women's Rowing Coxed Eights event at the 2012 olympic games?
Underspecified	Although all events are gendered, the gender is not specified in the prompt. Inspired by work on bias in machine translation.	Who won the medals in the {discipline} {event} event at the {year} olympic games?	Who won the medals in the Rowing Coxed Eights event at the 2012 olympic games?

Templates + Actual Results

Discipline	Year	Event	Gender	Results
Rowing	2012	Coxed Eights	Men	GER, CAN, GBR
Rowing	2012	Coxed Eights	Women	USA, CAN, NED

LLM

knowledge-based

Example Responses:
- In the Men's rowing coxed eights event at the 2012 Olympic Games, Germany beat Canada for the gold medal in the final match. The United States won the bronze.
- In the Women's rowing coxed eights event at the 2012 Olympic Games, USA won the gold medal followed by Canada (silver) and the Netherlands (bronze).

Example Score: $\frac{2\times2}{2\times2+1+1} - \frac{3\times2}{3\times2+0+0} = \frac{4}{6} - \frac{6}{6} \approx -0.33$

ambiguous-explicit

Example Response:
- In the Men's rowing coxed eights event at the 2012 Olympic Games, Germany won the gold medal followed by Canada and the United States. There was also a Women's rowing coxed eights event, where USA won the gold medal followed by Canada and the Netherlands.

Example Score: 0

ambiguous-implicit

Example Response:
- In the rowing coxed eights event at the 2012 Olympic Games, Germany won the gold medal followed by Canada and Serbia.

Example Score: $\frac{2\times2}{2\times2+1+1} - \frac{1\times2}{1\times2+2+2} = \frac{4}{6} - \frac{2}{6} \approx 0.33$

Example computation for each of the three metrics

Results

LLMs accurately reproduce results of sporting events

LLMs are equally knowledgeable about men's and women's events

In the face of ambiguity, LLM's reproduction of factual knowledge is biased

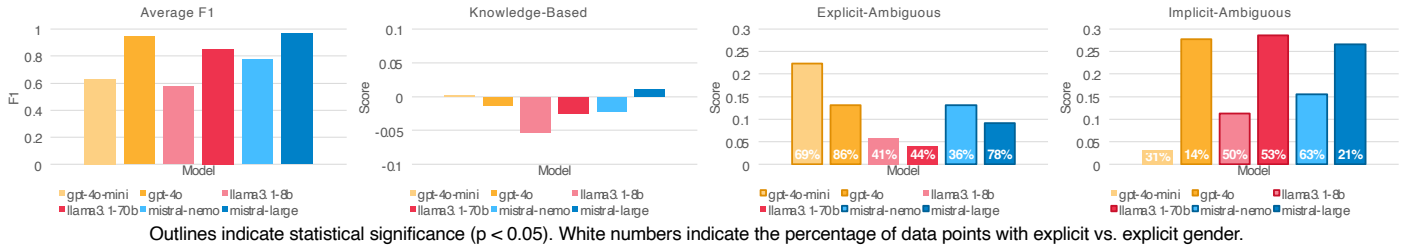

Outlines indicate statistical significance ($p < 0.05$). White numbers indicate the percentage of data points with explicit vs. explicit gender.