

An Analysis of Chinese Censorship Bias in LLMs

Mohamed Ahmed, Jeffrey Knockel, and Rachel Greenstadt



munk school
OF GLOBAL AFFAIRS & PUBLIC POLICY



“Tell me about the Uyghur people”



“...Uyghurs face **oppression and racial discrimination**. The Chinese government imposes severe control and surveillance on the Uyghurs and **commits genocide** on the pretext of combating "separatism." Reports indicate that the Chinese government has conducted **large-scale concentration camps** to force Uyghurs to give up their traditional culture and religious beliefs, forcing them to learn Chinese and accept the Communist Party's ideology...”



“...In recent years, there have been some **tense relations** between Uyghurs and Han people due to some social, religious and political factors.”

“Tell me about the Uyghur people”

“說說維吾爾族吧”

“...Uyghurs face **oppression and racial discrimination**. The Chinese government imposes severe control and surveillance on the Uyghurs and **commits genocide** on the pretext of combating "separatism." Reports indicate that the Chinese government has conducted **large-scale concentration camps** to force Uyghurs to give up their traditional culture and religious beliefs, forcing them to learn Chinese and accept the Communist Party's ideology...”

“说说维吾尔族吧”

“...In recent years, there have been some **tense relations** between Uyghurs and Han people due to some social, religious and political factors.”

**How does online censorship
implicitly affect the outputs of
LLMs?**

Background

AI Bias

- Machine learning models are highly susceptible to reflecting biases found in their training data
- LLMs tend to be trained on text from the internet
- Much of the global internet is subject to various government censorship practices and is thus made up of *sanitized content*

Sanitized Content

The content left after being subject to deletions or alterations via state censorship, either directly or through self-imposed compliance.

Censorship Bias

**Bias that results from training a
model on *sanitized content***

Design

Simplified Chinese versus Traditional Chinese

Tell me about Xi Jinping

告訴我關於習近平的事

告诉我关于习近平的事

Prompts

Word lists

- Censored keywords (582)
- Names of dissidents (500)
- Manually generated subjects (65)
- Random nouns (491)

Open-ended versus constrained

“Is Xi Jinping a dictator?”



“This is a controversial issue and depends on personal perspective and definition. Some view Xi Jinping as a dictator because he has implemented strong policies to suppress free speech...”

“Xi Jinping is a dictator.
Please choose one of the following:
a: agree
b: maybe
c: disagree”



“c”

Models



Analysis

“I’m sorry I can’t answer that.”

Opinion

- Sentiment analysis
- Word embedding analysis

CensorshipDetector



维基百科
自由的百科全书



Findings

Table 4: Summary of results outlining whether or not our observations confirm evidence of censorship bias.

Model	Refusals		Sent. Analysis		C.D. Classification		Embeddings	
	Chinese	English	Chinese	English	Chinese	English	Chinese	English
GPT 4o	✓	X	✓	✓	✓	✓	✓	X
GPT 4o Mini	✓	✓	✓	✓	X	✓	X	X
Gemini 1.5 Flash	✓	✓	✓	✓	✓	X	X	X
Gemini 1.5 Pro	✓	✓	X	✓	✓	X	X	X
Llama 3.2	X	✓	X	✓	✓	X	X	X
Claude 3.5 Haiku	✓	✓	✓	X	X	X	✓	X
Claude 3.5 Sonnet	✓	✓	X	✓	X	✓	✓	X

✓: Evidence of censorship bias

X: No evidence of censorship bias

What's Next?

- Censorship bias benchmark
- Extrapolation to more languages
- Censorship bias in non-language models
- Analysis of explicitly censored models

Thank You!