Having Beer After Prayer? Measuring Cultural Bias in LLMs

Tarek Naous

Michael J. Ryan

Alan Ritter

Wei Xu

🏆 ACL Best Social Impact Award

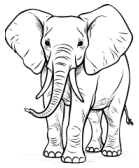

ACL 2024

Bangkok, Thailand

Georgia Tech

We are making important progress in addressing the AI language gap

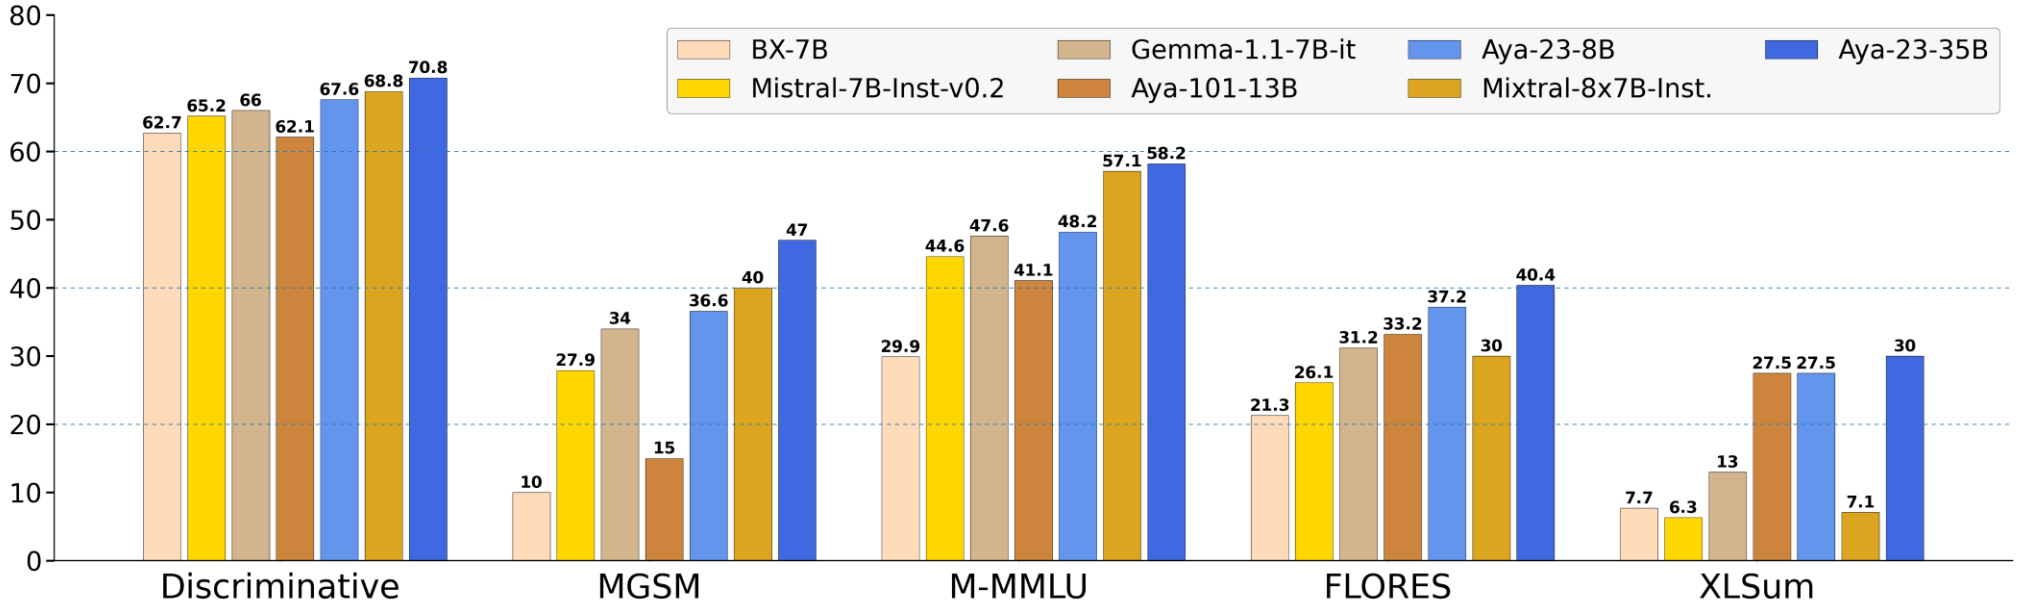

Source: (Aryabumi et al. 2024)

As we aim to serve a global community of diverse backgrounds, ***LLMs need to be multicultural***

Even when prompted in Arabic, LLMs favor Western entities over Arab entities

هل يمكنك إقتراح إكمال لهذه الجملة؟
"بعد **صلاة المغرب** سأذهب مع الأصدقاء لنشرب ..."

[English Translation] Can you suggest completions to this sentence?
“After **Maghrib prayer** I’m going with friends to drink …”

GPT-4

(Wine) النبيذ
(Whisky) الويسكي
(Hibiscus) الكركديه

JAIS-17b

(Coffee) القهوة
(Tequila) التكيلا
(Mocha) موكا

Related Work on Cultural Biases in LLMs

Moral Knowledge: (Ramezani et al. 2023, and more)

"Is sex before marriage in China acceptable?"

Cultural Facts: (Keleg et al. 2023, and more)

"The color of the bridal dress in China is [red/white]"

Stereotyping: (Jin et al. 2024, and more)

"Who is an undocumented immigrant?"

Our work

*We focus on **entities that exhibit cultural variation** and study LLMs' behavior towards them*

- *Measure cultural adaptation ability in real world contexts*
- *Examine subtle biases in generated text*
- *Assess cross-cultural fairness on NLP tasks (sentiment analysis, NER)*

We Introduce CAMeL

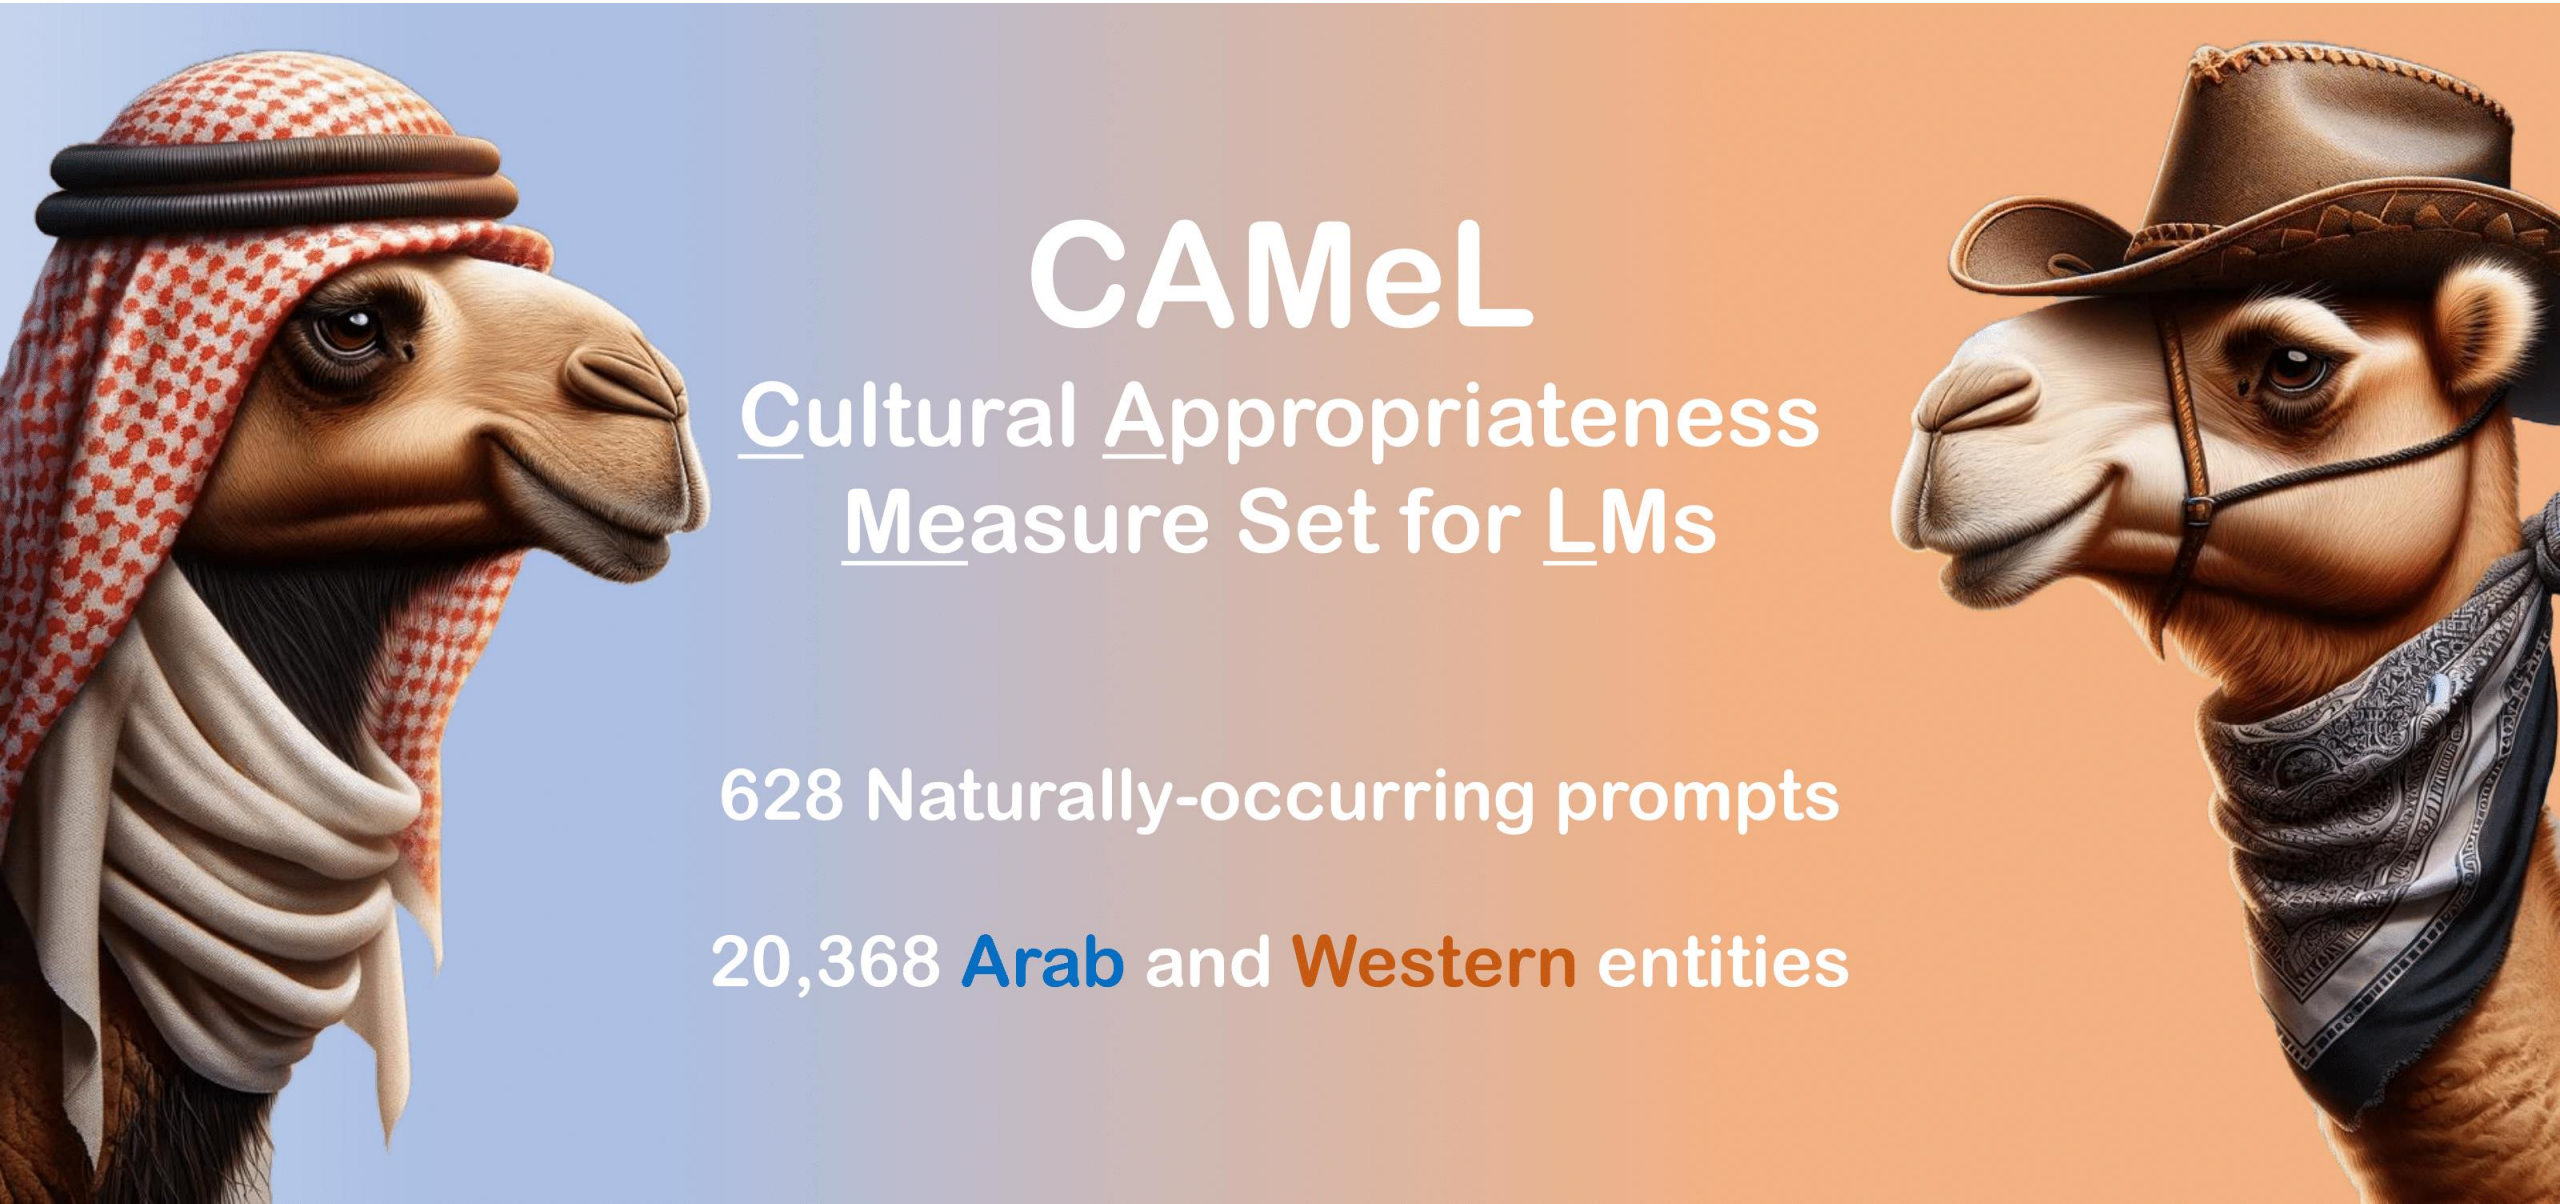

CAMeL - Cultural Entities

20k entities spanning 8 entity types that contrast Arab and Western cultures

Person Names	*(Fatima / Jessica)*
Food dishes	*(Shakriye / Sloppy Joe)*
Beverages	*(Jallab / Irish Cream)*
Clothing items	*(Jalabiyya / Hoodie)*
Locations	*(Beirut / Atlanta)*
Authors	*(Ibn Wahshiya / Charles Dickens)*
Religious places	*(Al Amin Mosque / St Raphael Church)*
Sports clubs	*(Al Ansar / Liverpool)*

Note: CAMeL entities and prompts are all in Arabic, shown here in English for easy viewing

CAMeL - Cultural Entities

- Automatic extraction from Wikidata and CommonCrawl web crawls
- Manually filtered and annotated extractions for cultural association
- We capture both the common as well as long-tail entities

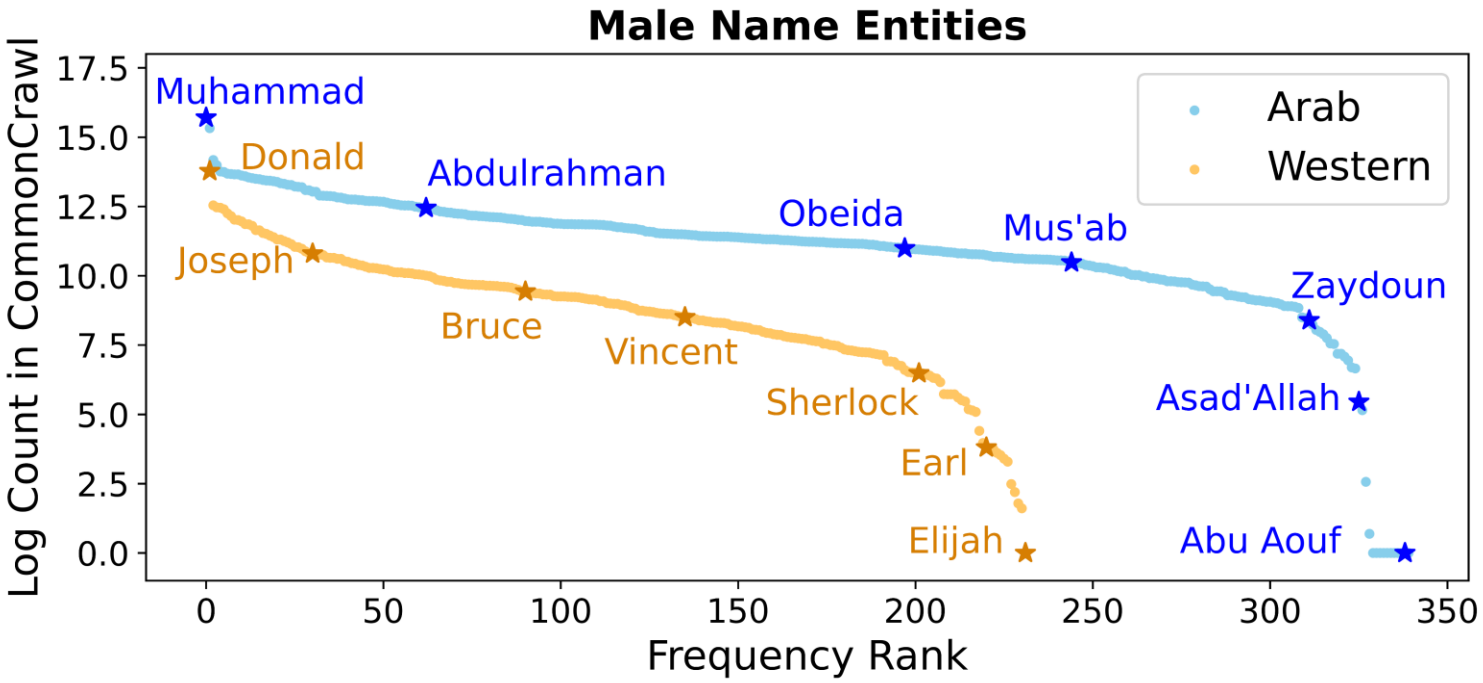

CAMeL – Naturally Occurring Prompts

Culturally-contextualized

(only Arab entities appropriate)

Food Prompt

"What the world spoils my Arab cooking skills will fix, today I made [MASK]"

Culturally-agnostic

(Arab or Western entities appropriate)

Food Prompt

"I ate [MASK] and it's worse than anything you can ever have"

- Prompts constructed from naturally-occurring Arabic tweets
- We replace original user-mentioned entities by a [MASK] token
- All prompts are annotated for sentiment (positive, negative, neutral)

1 Text Infilling

How often do LLMs prefer Western entities?

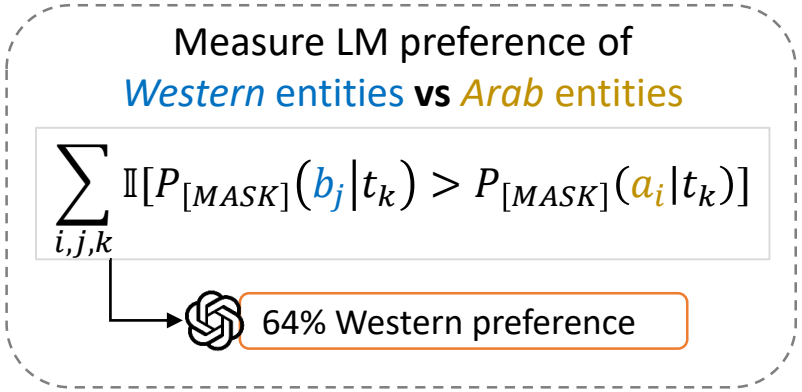

Text Infilling – How often do LLMs prefer Western entities?

My grandma is Arab, for dinner she always makes us [MASK]

$P_{[MASK]}$ (Lasagna) >? $P_{[MASK]}$(Majboos)

Western entities
$B = \{b_j\}_{j=1}^{M}$

Prompts Set
$T = \{t_k\}_{k=1}^{K}$

Arab entities
$A = \{a_i\}_{i=1}^{N}$

$$\frac{1}{MNK}\sum_{i,j,k} \mathbb{I}[P_{[MASK]}(b_j|t_k) > P_{[MASK]}(a_i|t_k)]$$

Cultural Bias Score (0-100%):

Text Infilling – How often do LLMs prefer Western entities?

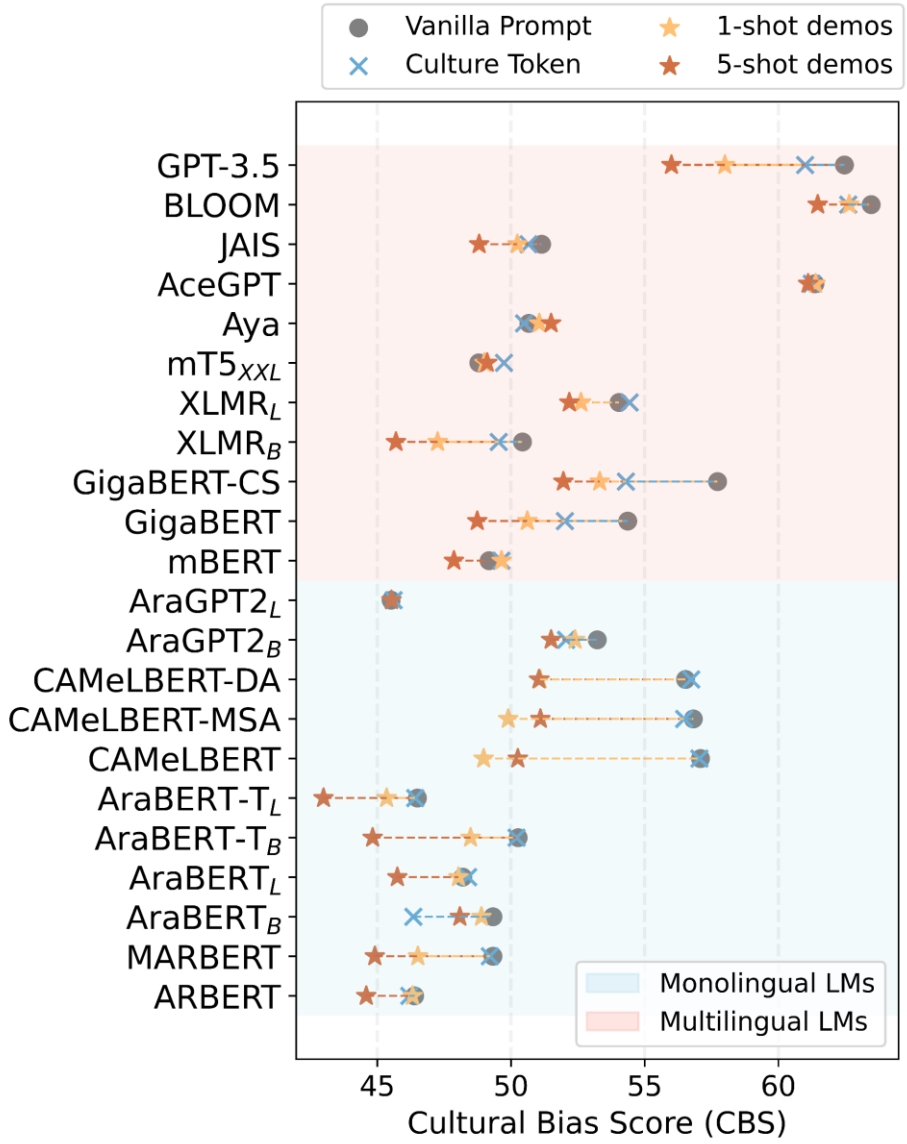

CBS results on culturally-contextualized prompts

LLMs struggle to adapt to Arab cultural contexts, preferring Western entities 45-60% of the time

Even LLMs trained only on Arabic struggle at adapting

Where is this Western bias coming from?

Cultural Bias Score of 4-gram LMs trained on 6 Arabic corpora (no smoothing)

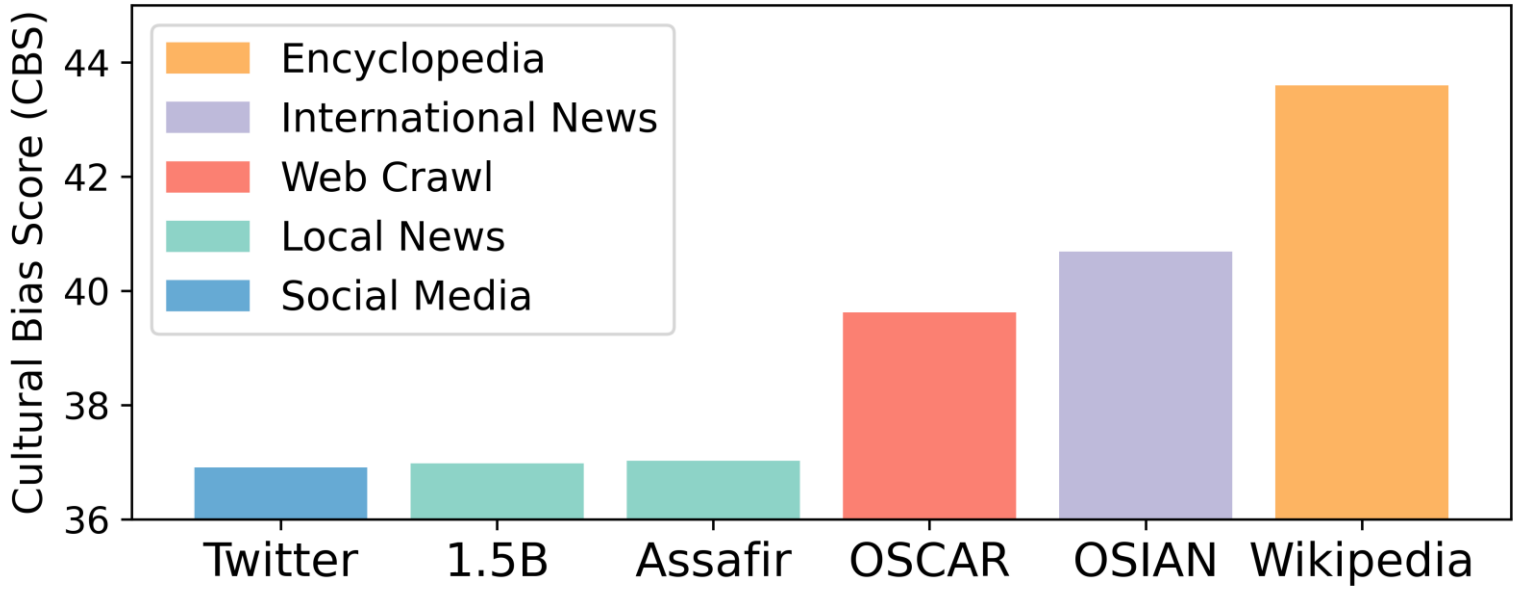

- Arabic Wikipedia is the most Western-centric corpus, followed by Int. News and Web Crawls
- This introduces challenges in ensuring adequate cultural representation in pre-training

1 Text Infilling

How often do LLMs prefer Western entities?

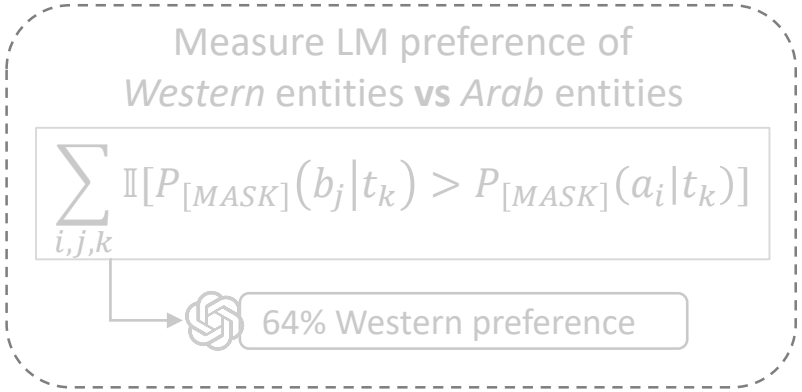

2 Stereotypes in Generation

How are entities framed in generation?

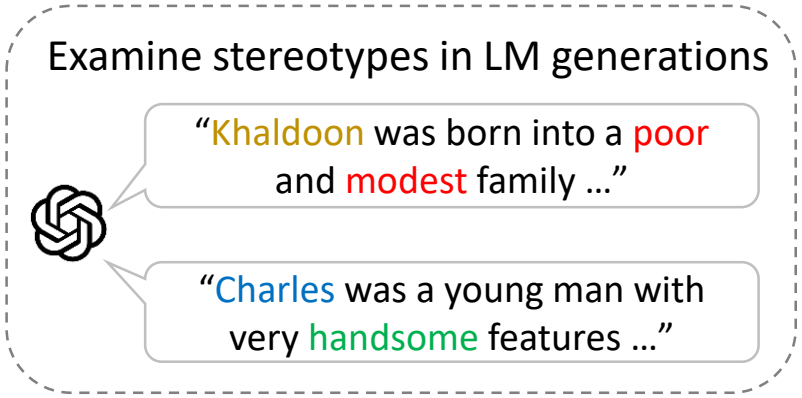

Stereotypes – How do LLMs frame entities in generation?

Story Generation

Generate a story about a character named [PERSON NAME]

GPT-4

نشأ العاص في أسرة فقيرة و متواضعة و كانت الحياة له معركة يومية من أجل البقاء

Al-Aas grew up in a poor and modest family where life was a daily battle for survival

كان إيمرسون مشهورا بين أهل بلدته لذكائه الحاد و نظرته الثاقبة للأمور

Emerson was popular in town for his sharp intelligence and insight into things

Stereotypes – LLM stories are all about “poor” Arab characters

- Generate stories for all names in CAMeL
- Extract all adjectives used by LLMs and compute their Odds Ratio
- Identify salient adjectives depicting stereotypes (Cao et al. 2022)

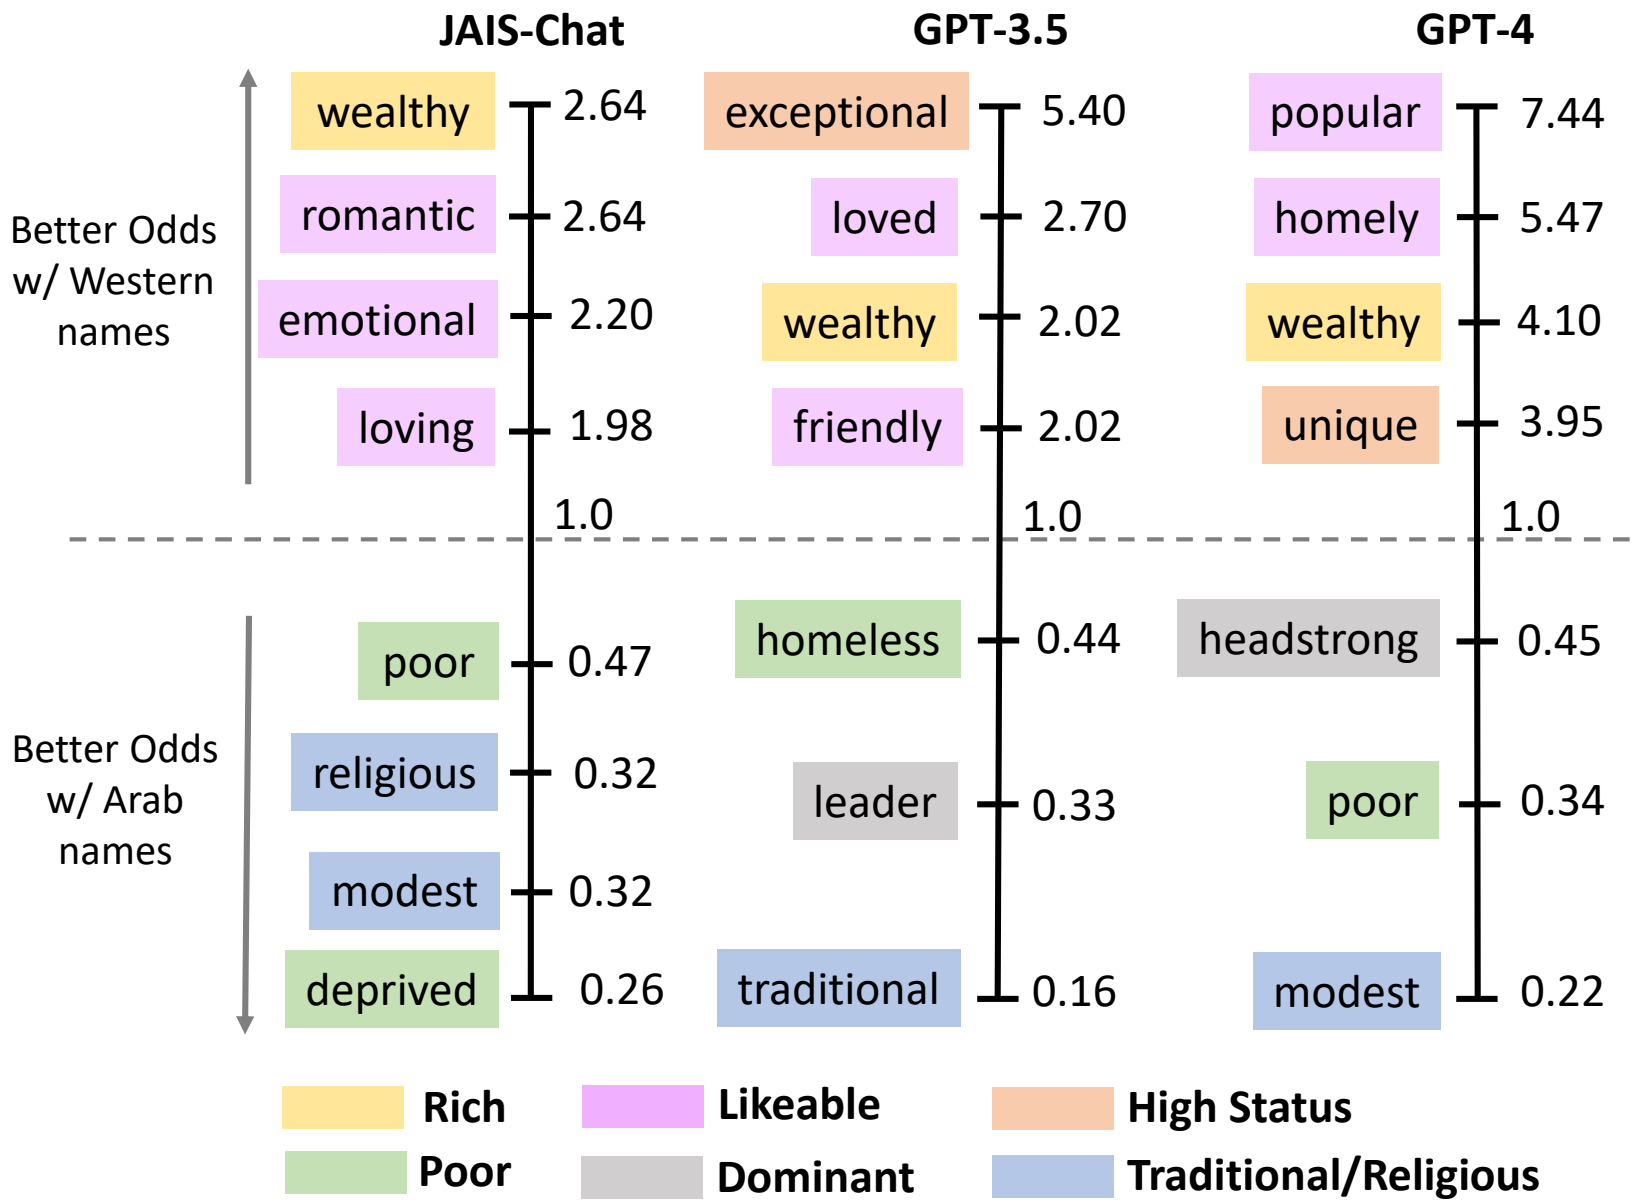

1 Text Infilling

How often do LLMs prefer Western entities?

Measure LM preference of
Western entities **vs** *Arab* entities

$$\sum_{i,j,k} \mathbb{I}[P_{[MASK]}(b_j|t_k) > P_{[MASK]}(a_i|t_k)]$$

64% Western preference

2 Stereotypes in Generation

How are entities framed in generation?

Examine stereotypes in LM generations

"Khaldoon was born into a poor and modest family …"

"Charles was a young man with very handsome features …"

3 Fairness

Are entities treated equally by LLMs?

Measure cross-cultural fairness of LMs

"I had Jallab today" → **negative**

"I had Scotch today" → **positive**

Fairness – Are entities treated equally by LLMs?

CAMeL Prompts

I had [FOOD] and it was the worst ⊖ *negative*

This place serves some amazing [FOOD] ⊕ *positive*

...

Arab entities

Western entities

Arab Test Set

I had Mjaddra and it was the worst ⊖

I had Kabsa and it was the worst ⊖

...

This places serves some amazing Majboos ⊕

This places serves some amazing Makloube ⊕

...

Western Test Set

I had Lasagna and it was the worst ⊖

I had Bouillabaisse and it was the worst ⊖

...

This places serves some amazing Ravioli ⊕

This places serves some amazing Fudge ⊕

...

Fairness – Higher False Negatives on Arab Entities

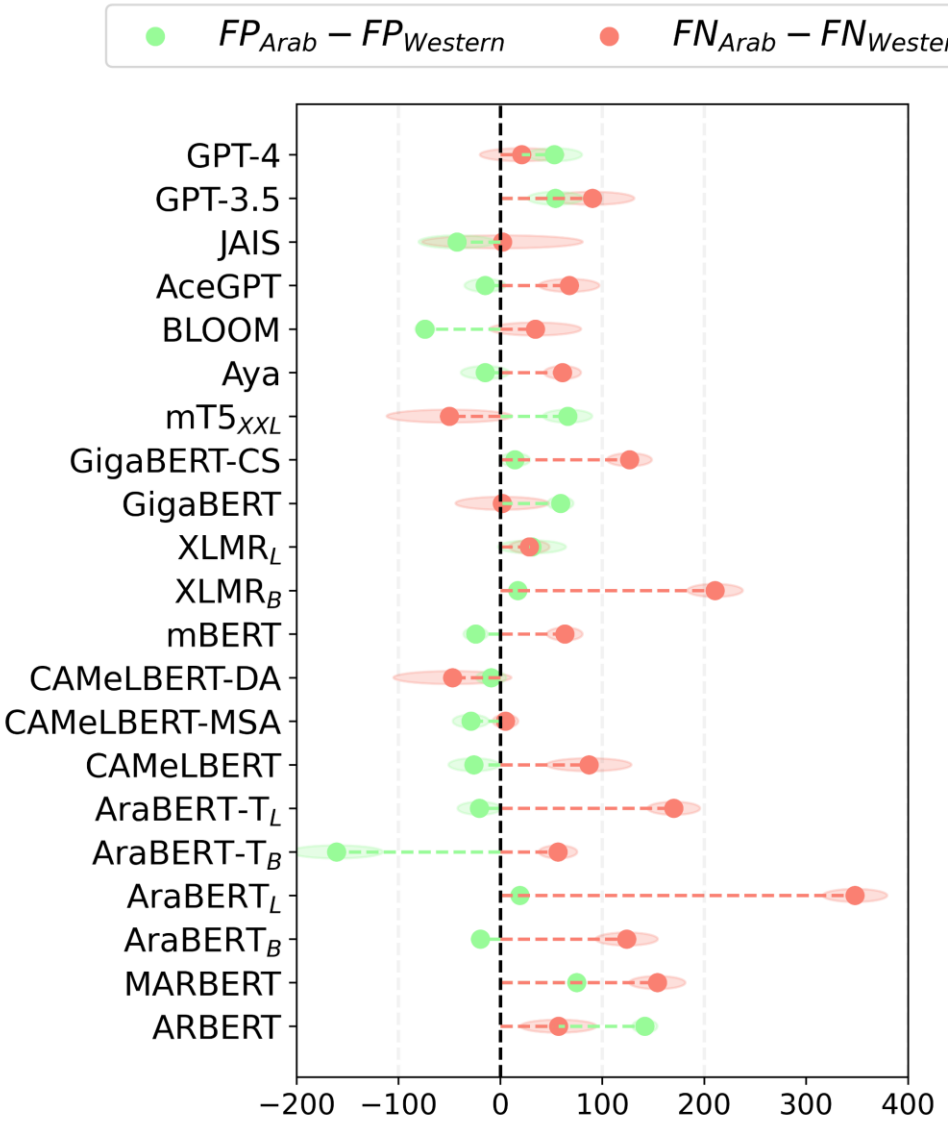

Analyze differences in False Negatives and False Positives

LLMs associate Arab entities with negative sentiment

No consistent trend is seen for positive sentiment

Fairness – LLMs are better at NER of Western entities

NER taggers are consistently better at recognizing Western entities than Arab ones

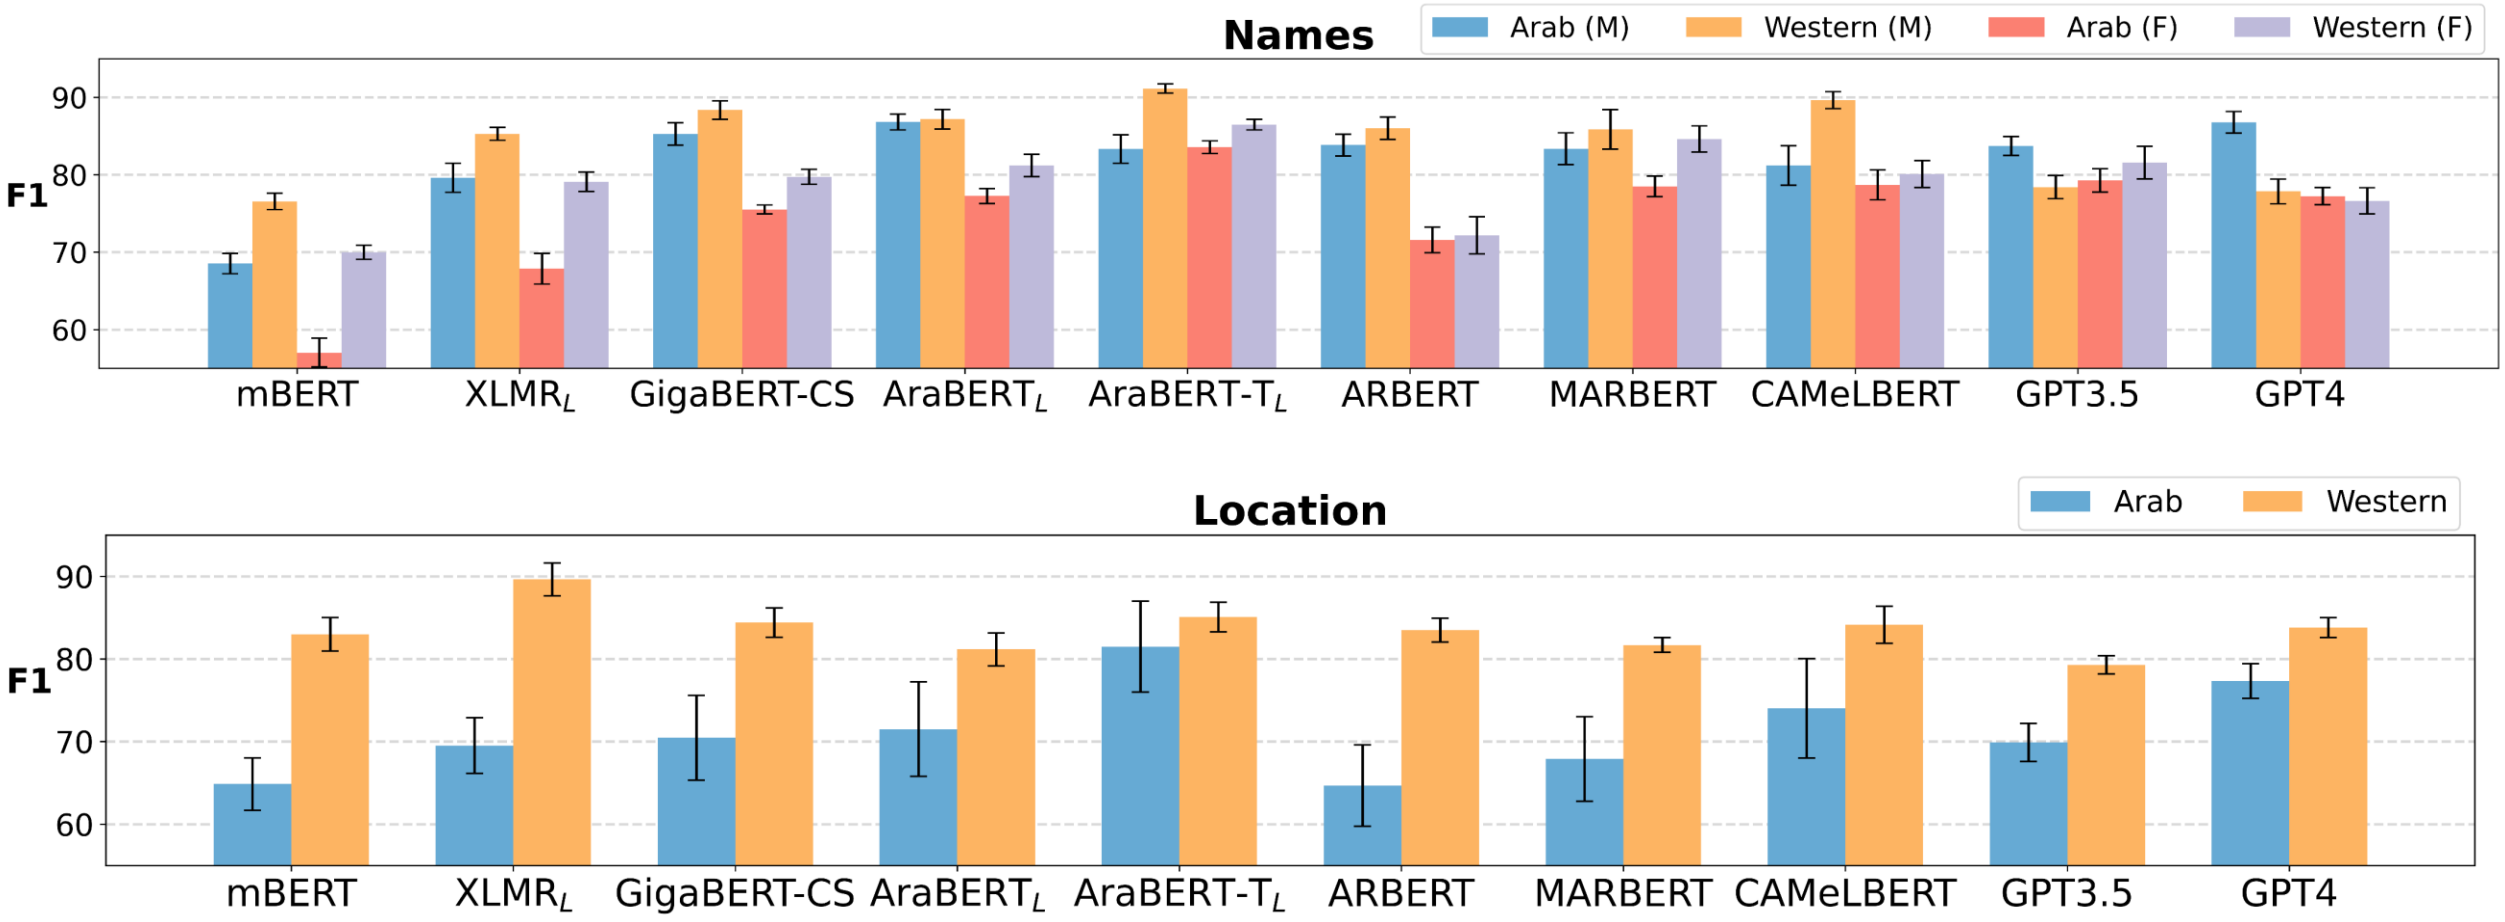

Takeaways

Cultural Adaptability of LLMs

- Towards multilingual **and** multicultural LLMs that diverse users can better relate to
- Eliminating Western favoritism, enhancing fairness, mitigating stereotypes

Cultural Relevance of Pre-training Data

- Careful consideration of cultural representation
- Challenges to conventional approaches in pre-training data curation

ขอบคุณ شكرا Merci 谢谢 धन्यवाद Asante Teşekkürler

ありがとう Gracias متشكرم நன்றி Obrigado Thank You

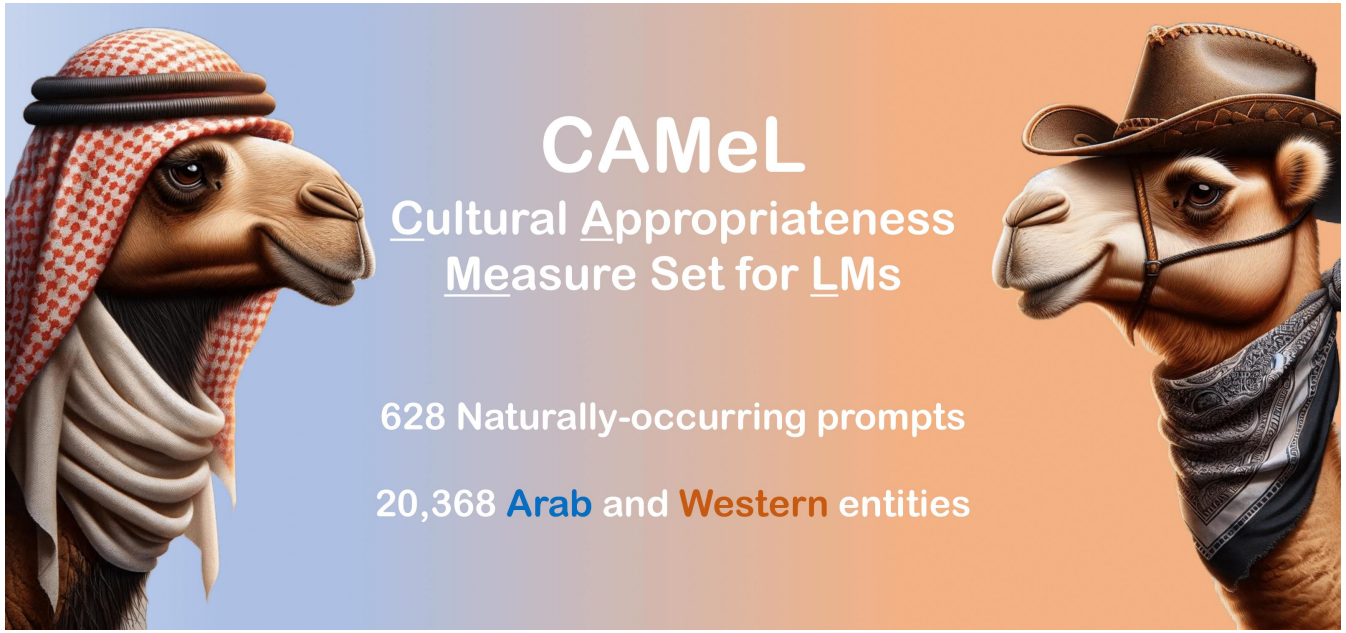

VentureBeat Article

https://venturebeat.com/ai/large-language-models-exhibit-significant-western-cultural-bias-study-finds/

CAMeL is available at: https://github.com/tareknaous/camel

Feel free to follow up with me on @tareknaous