

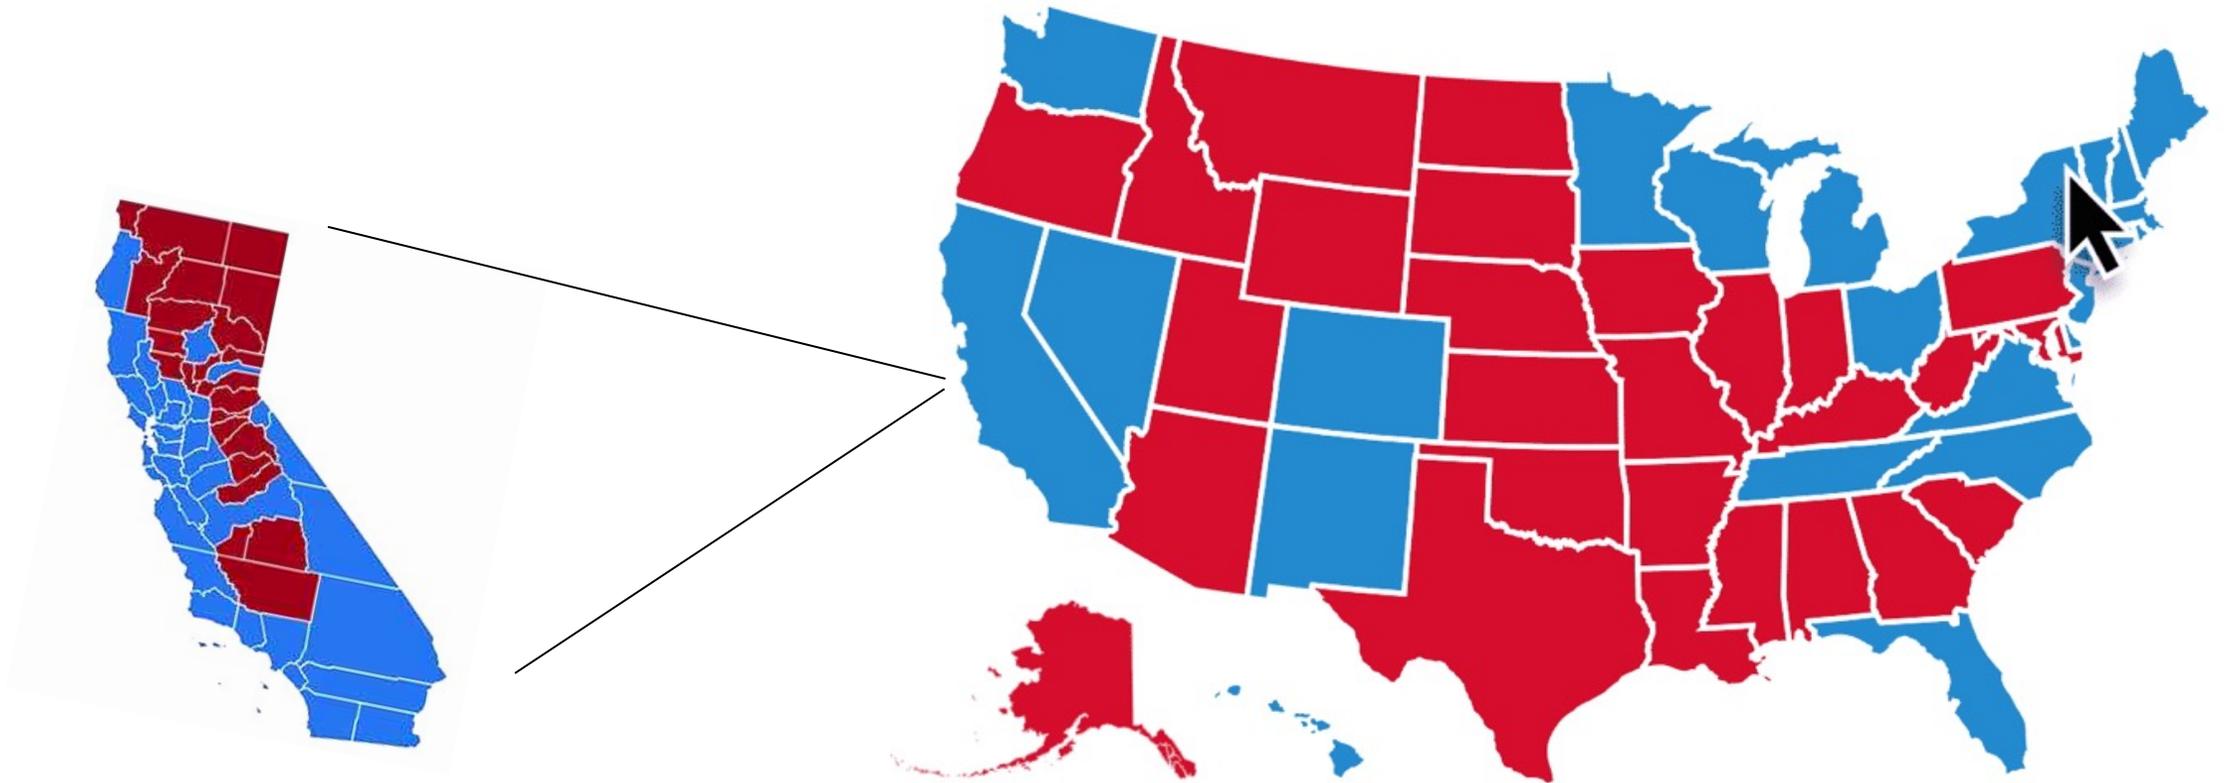
Building Knowledge-Guided Lexica to Model Cultural Variation



Shreya Havaldar, Salvatore Giorgi, Sunny Rai, Thomas Talhelm, Sharath Chandra Guntuku, Lyle Ungar

University of Pennsylvania

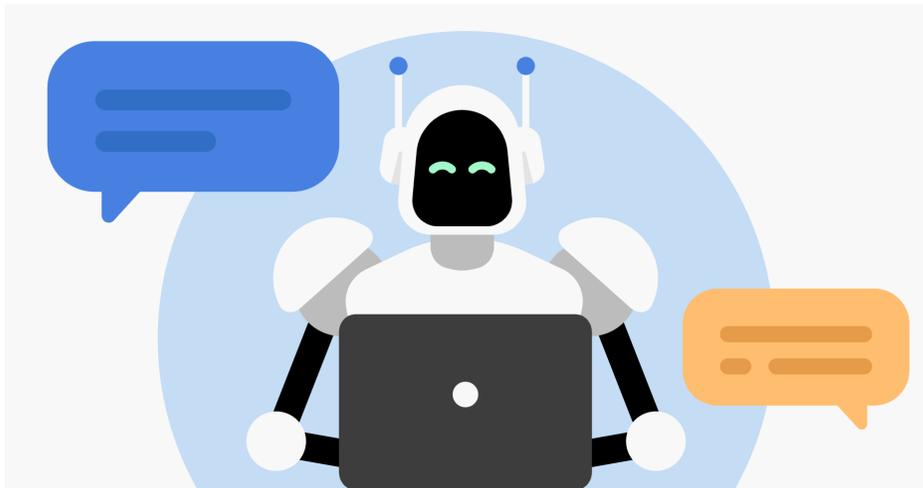
Culture is not homogenous! It has *variation*



2020 Presidential Election map
Trump (red) vs. Biden (blue)

Why do we want to measure regional variation in culture?

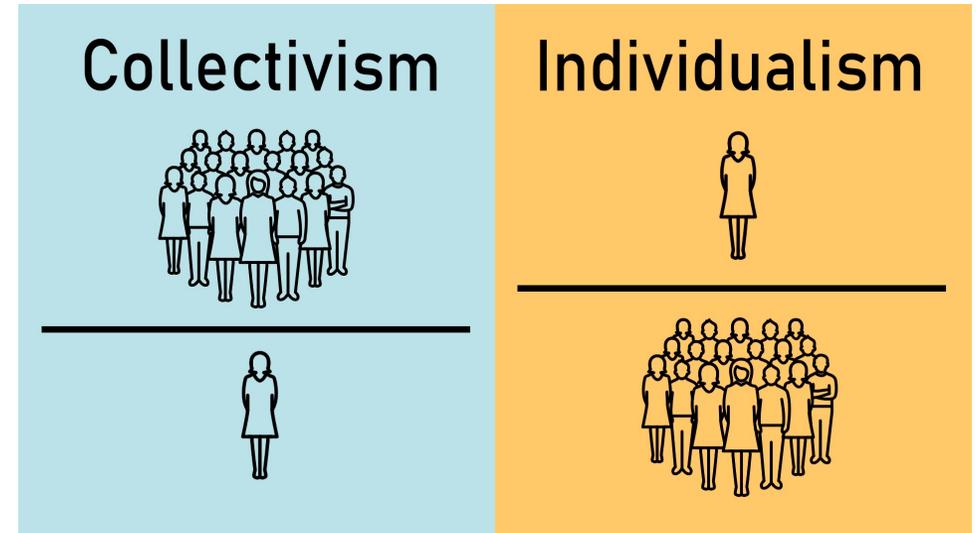
We can better **understand people** and their communities



We can **help LLMs adapt** to real-world cultural variation when generating language

Individualism and Collectivism

- We can tackle this problem by studying specific ***cultural constructs***
- An example construct is individualism and collectivism



Collectivism stresses the importance of the **community**, while individualism is focused on the rights and concerns of **each person**.

Are you more individualist or collectivist?

Do you agree with the following statements?

1. We should keep our aging parents with us at home
2. Children should be taught to place duty before pleasure
3. You should loan a good friend money with no questions asked

Are you more individualist or collectivist?

Do you agree with the following statements?

1. We should keep our aging parents with us at home
2. Children should be taught to place duty before pleasure
3. You should loan a good friend money with no questions asked

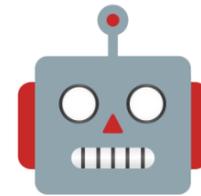
“Yes” → You are likely more **collectivist** than **individualist**

Your cultural background influences your language - *LLMs should reflect this.*

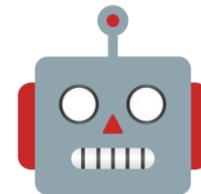


My elderly parents live at home with me.

That's great to hear! It is wonderful that you can care for them in their old age.



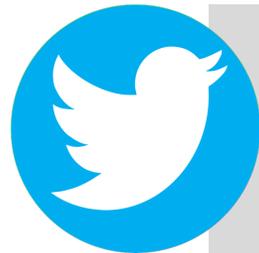
That sounds like it may be challenging. If you need any advice, feel free to ask!



Current methods of measuring cultural constructs are not scalable

1. Questionnaires/Surveys

- Too much overhead for researchers
- We don't want to use ***a small sample of people*** when measuring culture!



Idea: We can use peoples' language to measure cultural constructs

Current methods of measuring cultural constructs are not scalable

1. Labelling data at scale is expensive
 - Labeling 1 billion Tweets using GPT-4 would cost ~\$300,000
2. Model training at scale is expensive
3. Not all language is culturally relevant
 - How do we determine which Tweets indicate a cultural construct?



Measuring cultural variation is not a classic ML Problem!

Knowledge-guided lexica provide a *scalable* way to measure regional variation in culture

We generate sets of words (lexica) that characterize a cultural construct

- **No need to label data**
- **Computationally inexpensive**
- **Can utilize a massive dataset**
- **Inherently interpretable**



Using our lexical model, we can analyze ~1.5 billion Tweets from 6 million users in 20 minutes!

Traditional lexica have issues. We try to fix them!

1. Curated seed words from a domain expert
2. Embedding-based expansion
 - a) Improves **recall** – finds missing words that should be included
3. Lexica Purification
 - a) Improves **precision** – removes erroneous words that muddy results

We will apply this method to measure variation in individualism/collectivism across the United States

Curated seed words to measure individualism and collectivism

Humanity

Worldwide

Global

Equity

Cooperation

Identity

Guilt

Diversity

Individualism seed words

Duties

Responsibilities

Community

Sacrifice

Shame

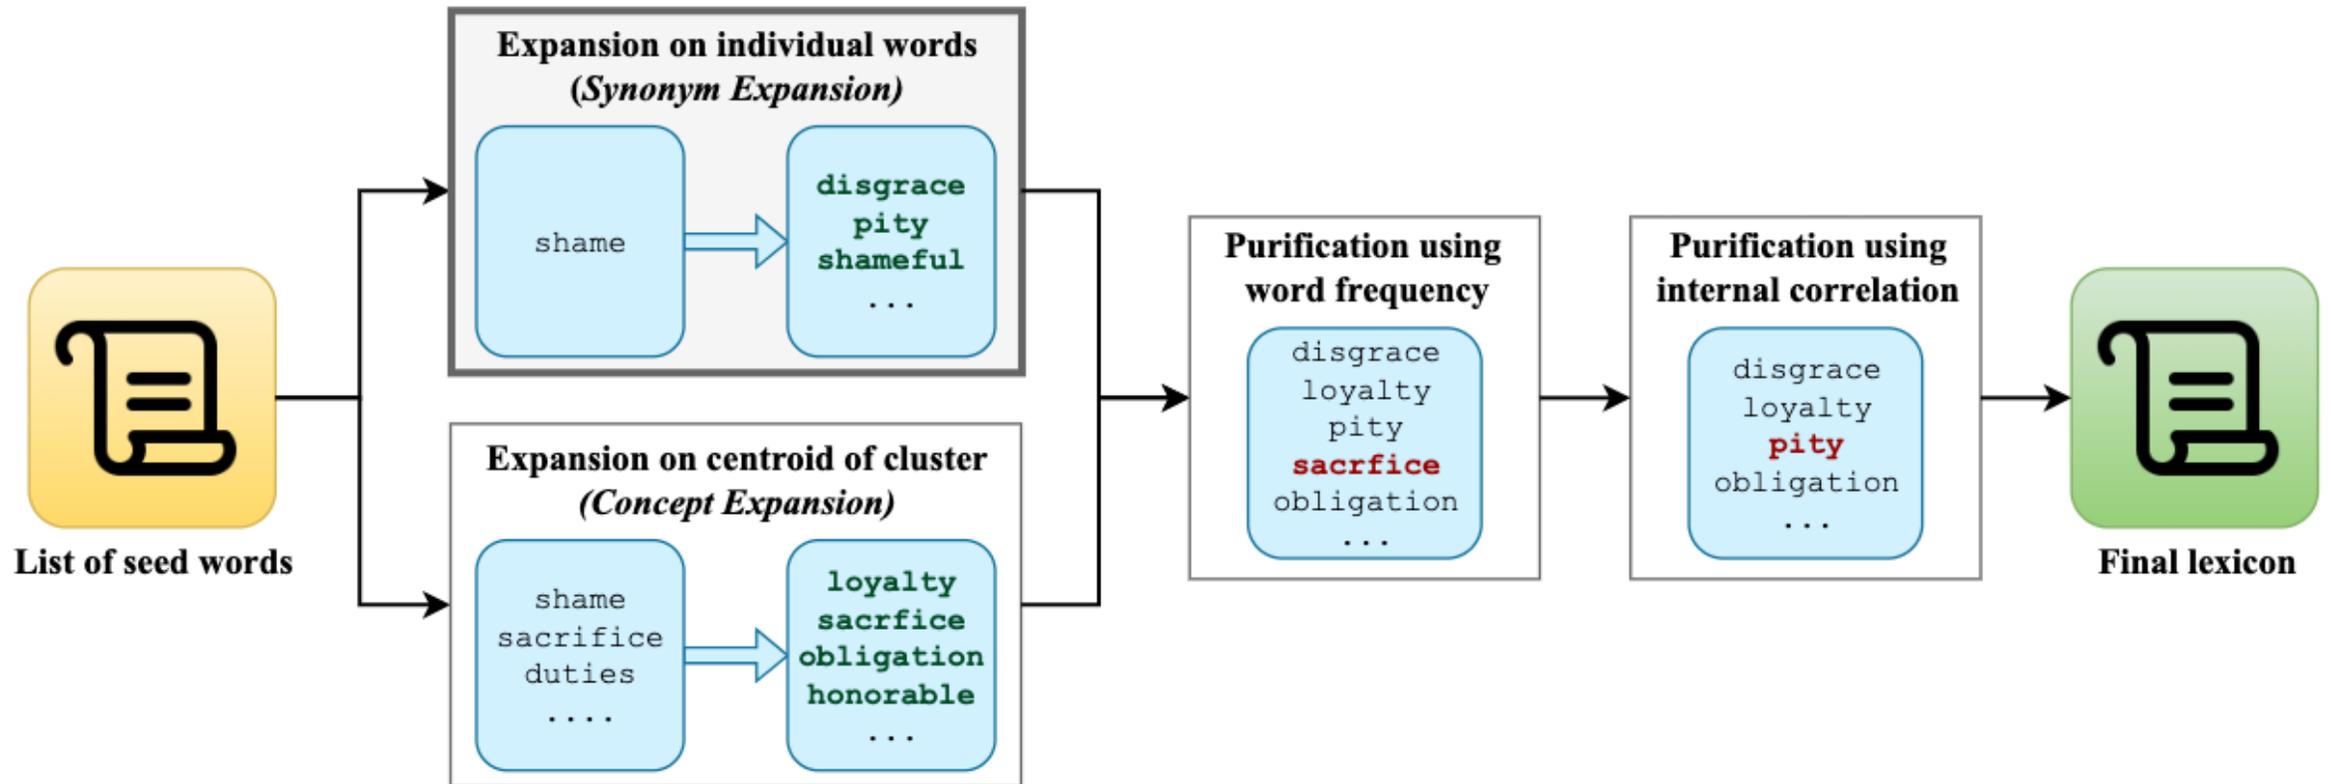
Rules

Honor

Obedience

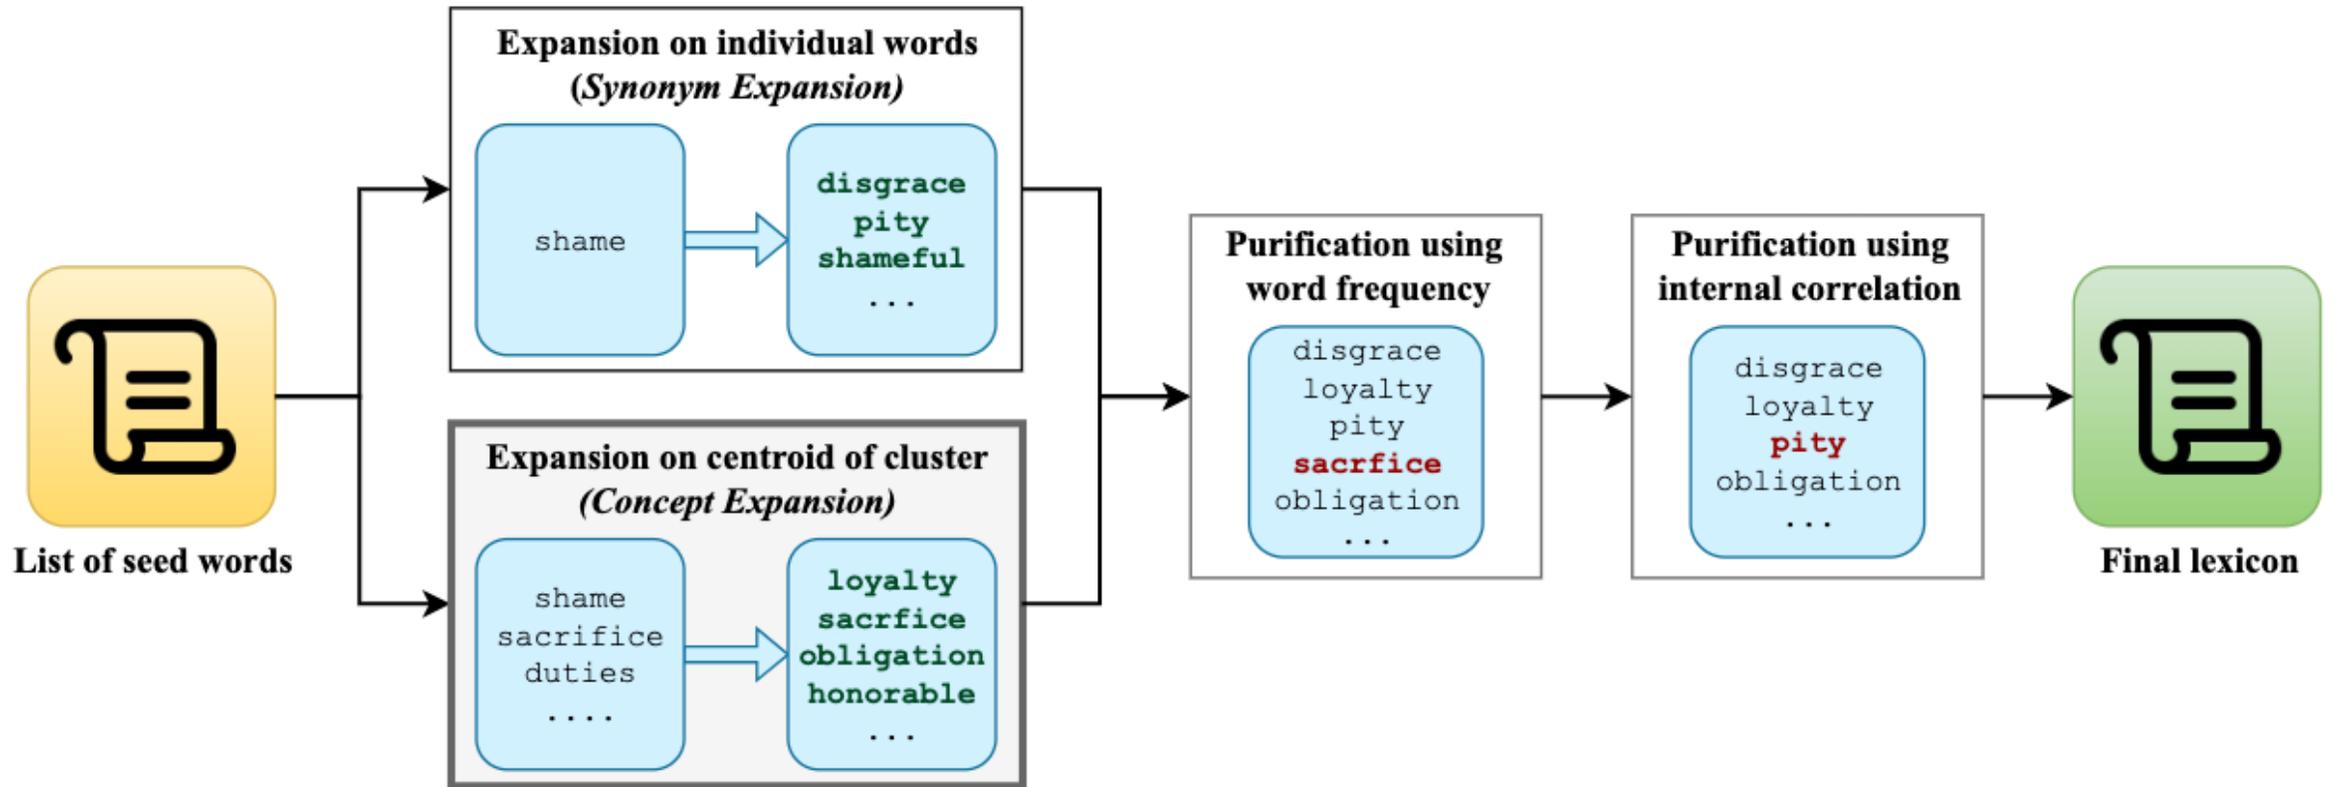
Collectivism seed words

Synonym Expansion



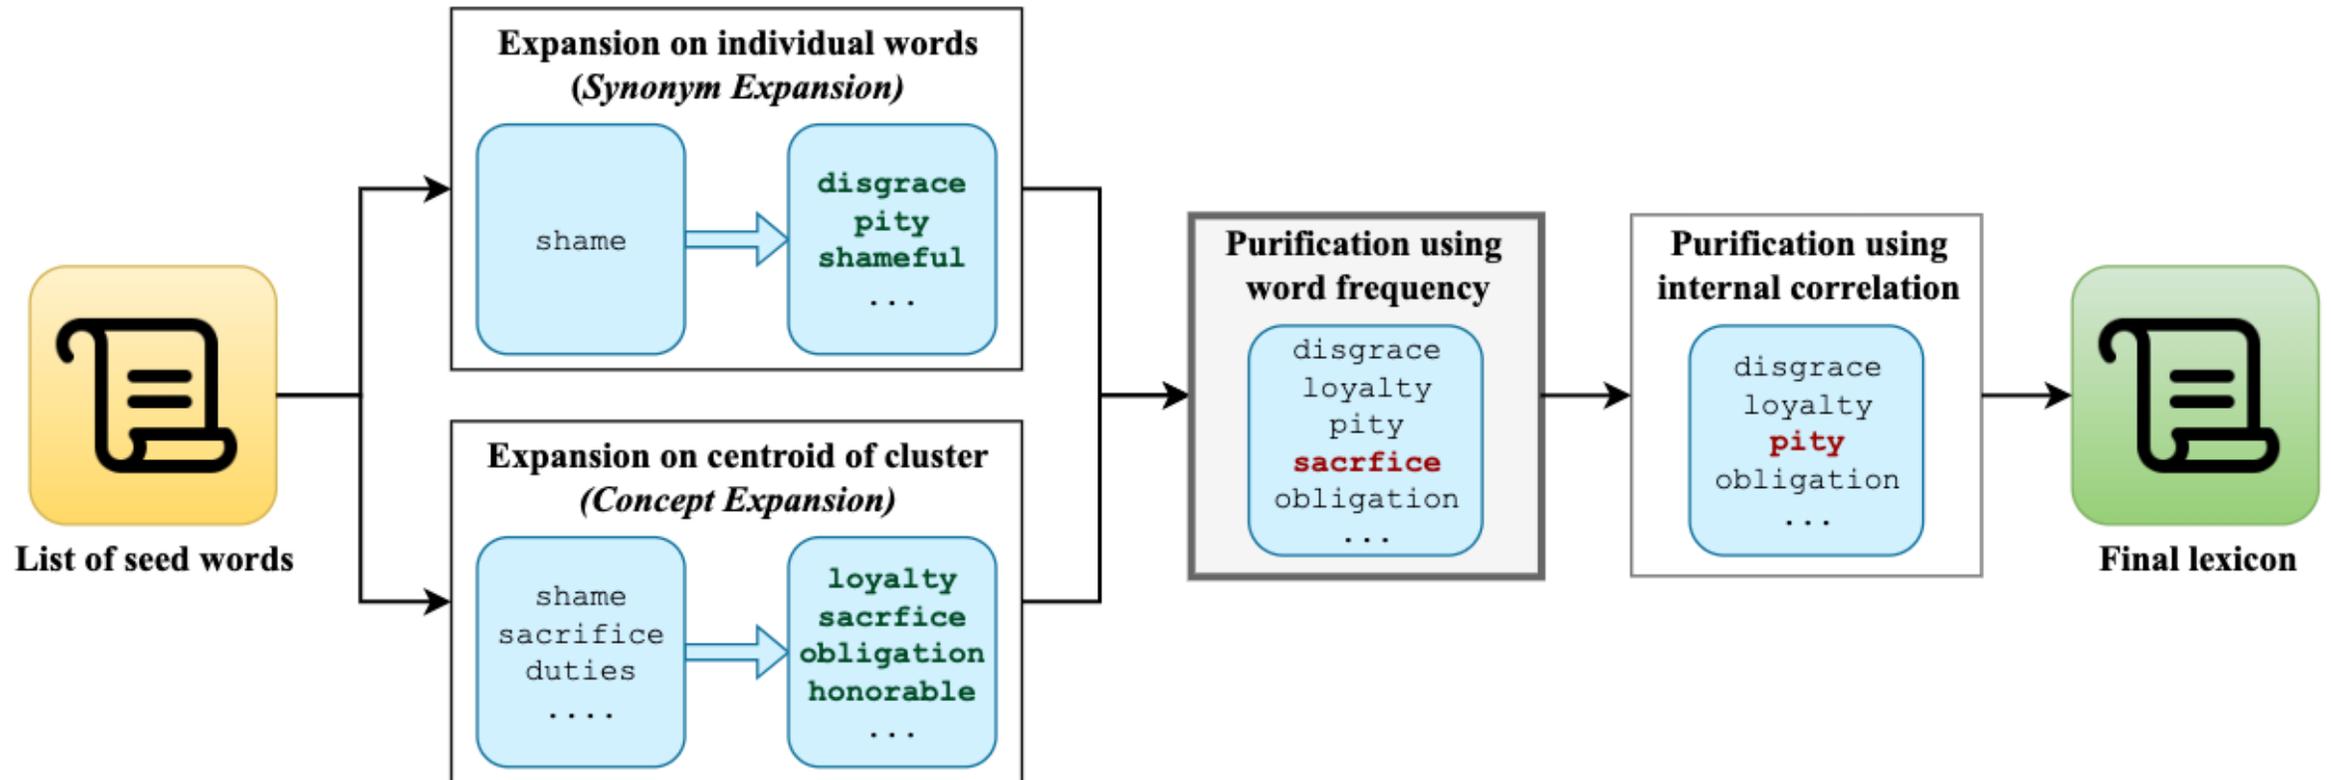
Find words closest to each individual word

Concept Expansion



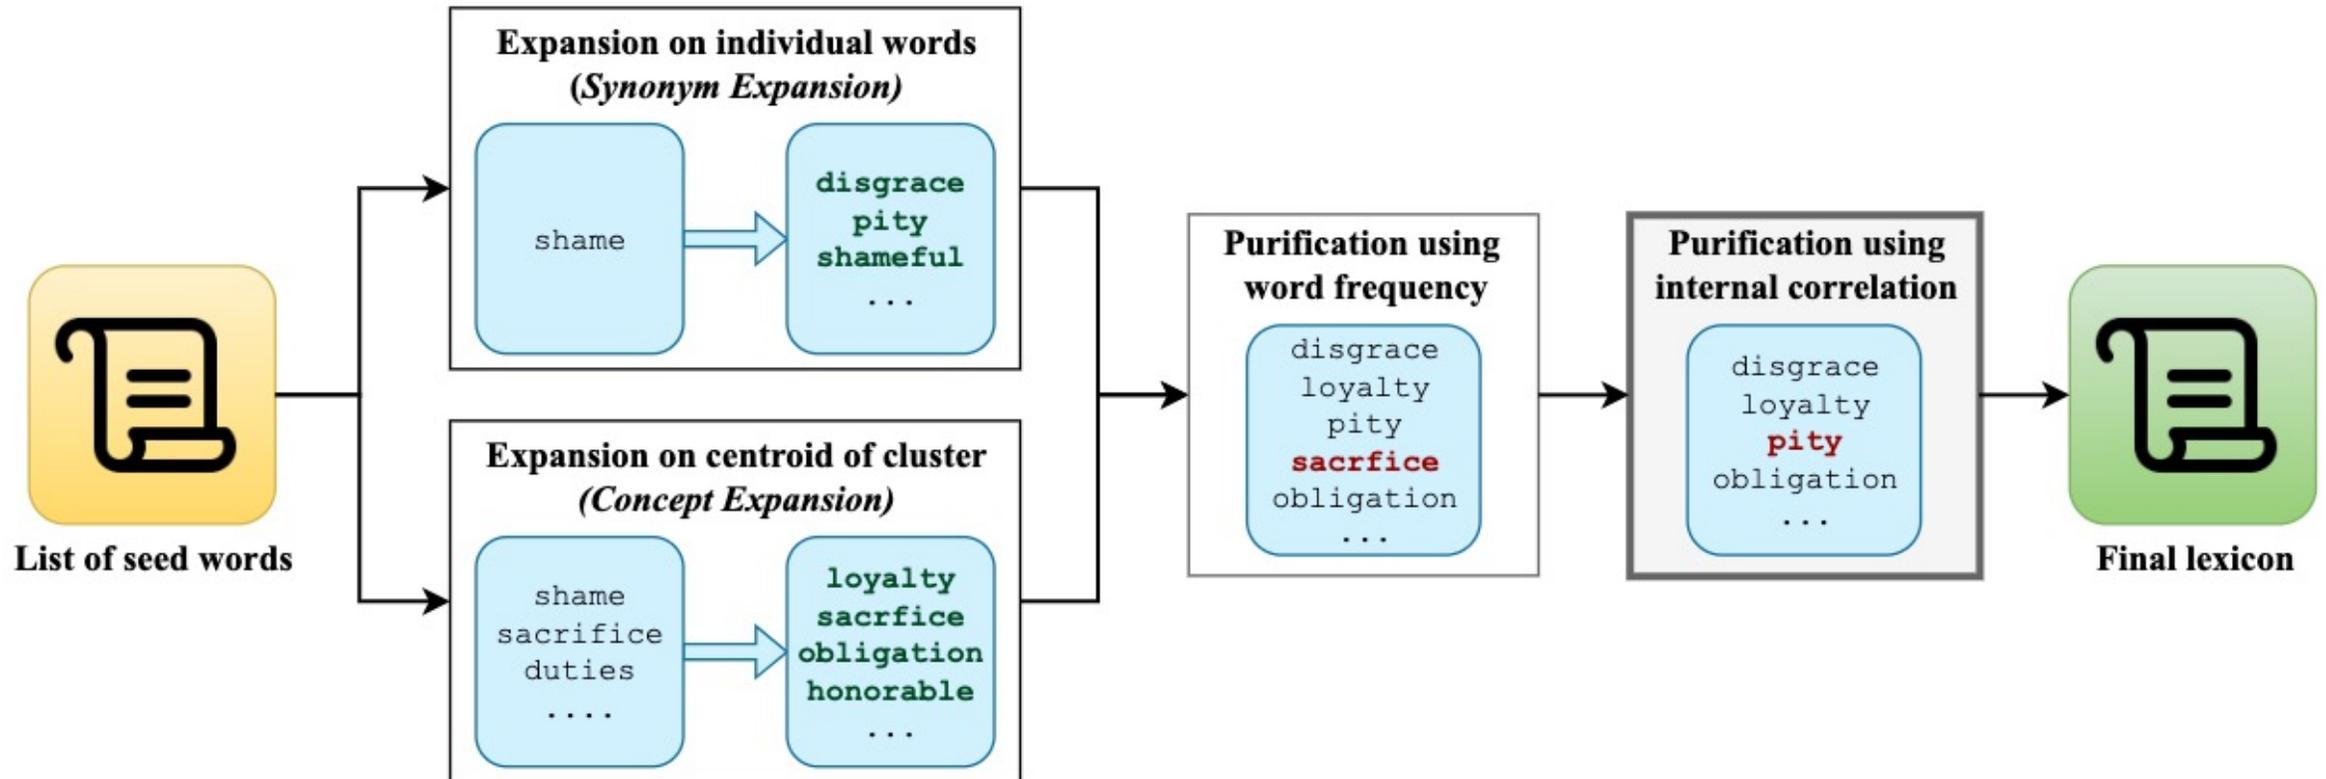
Find words closest to each category's centroid

Frequency-based purification



Remove rare words

Correlation-based purification



Remove words that do not co-occur with their category

10 seed words → ~200 word lexical model



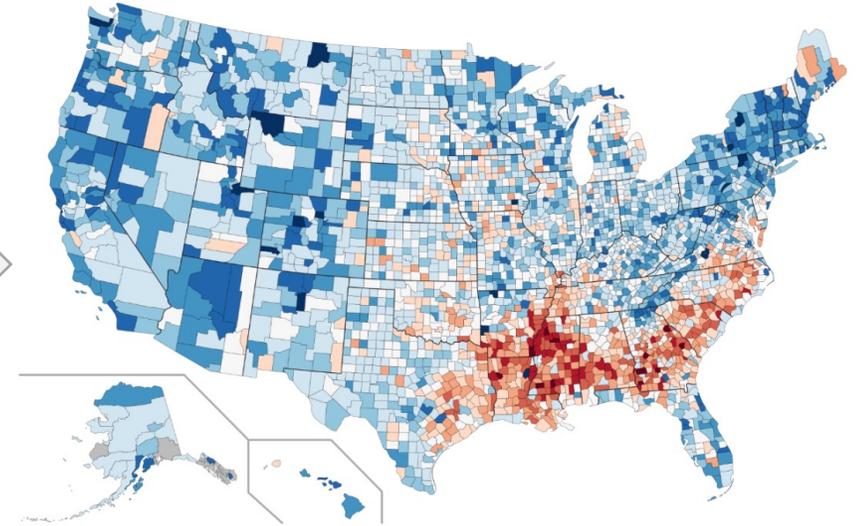
Individualism Lexica



Collectivism Lexica



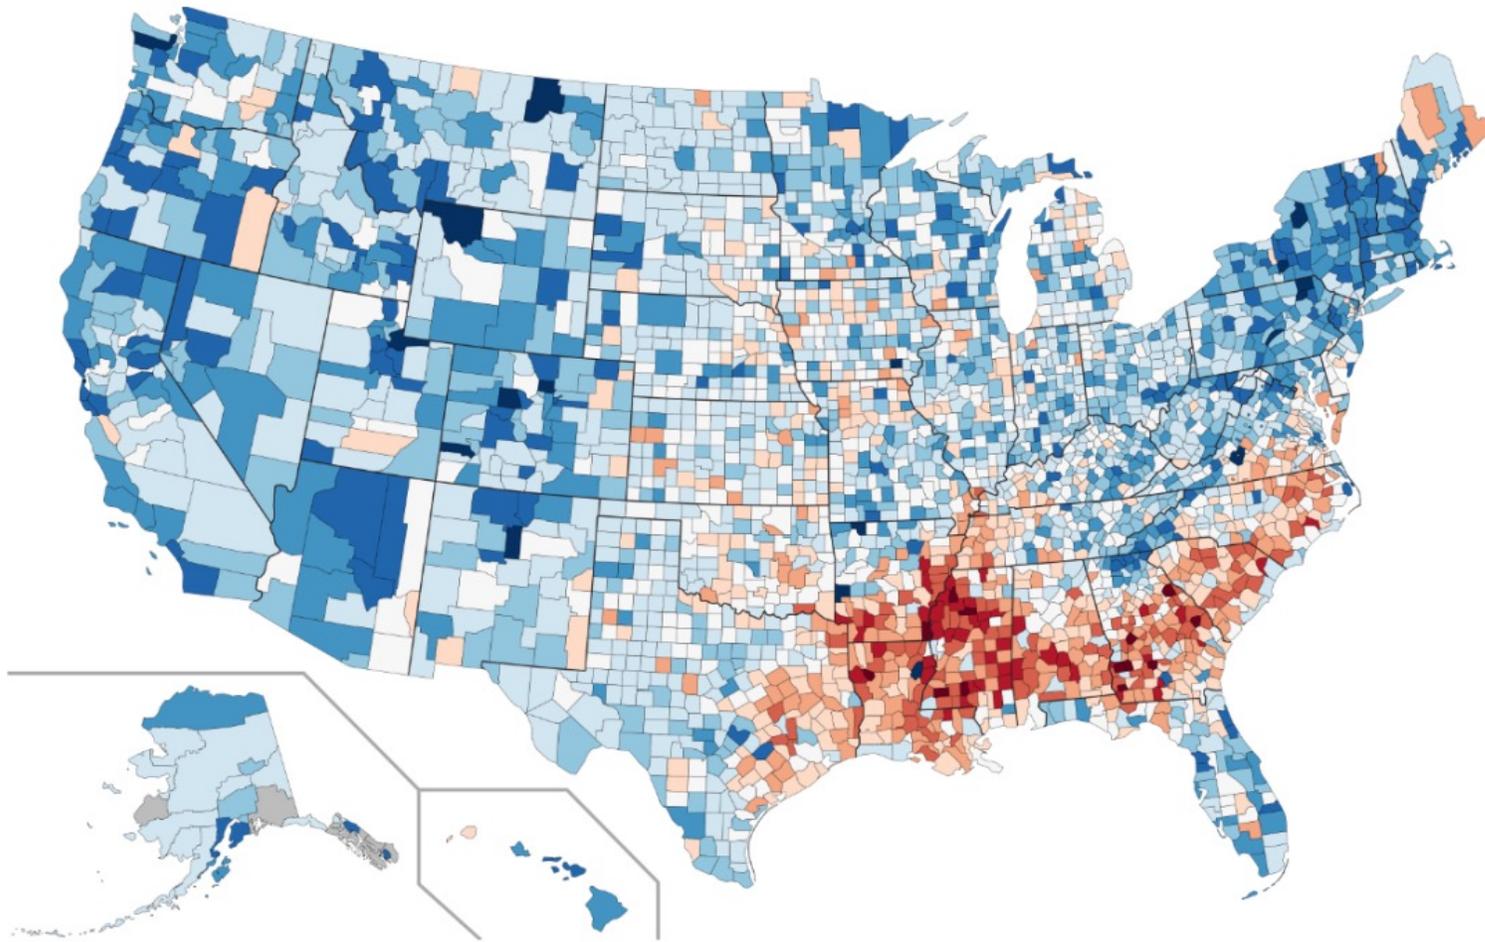
Knowledge-guided
lexical models



1.5 billion Tweets from
6 million geolocated
Twitter users

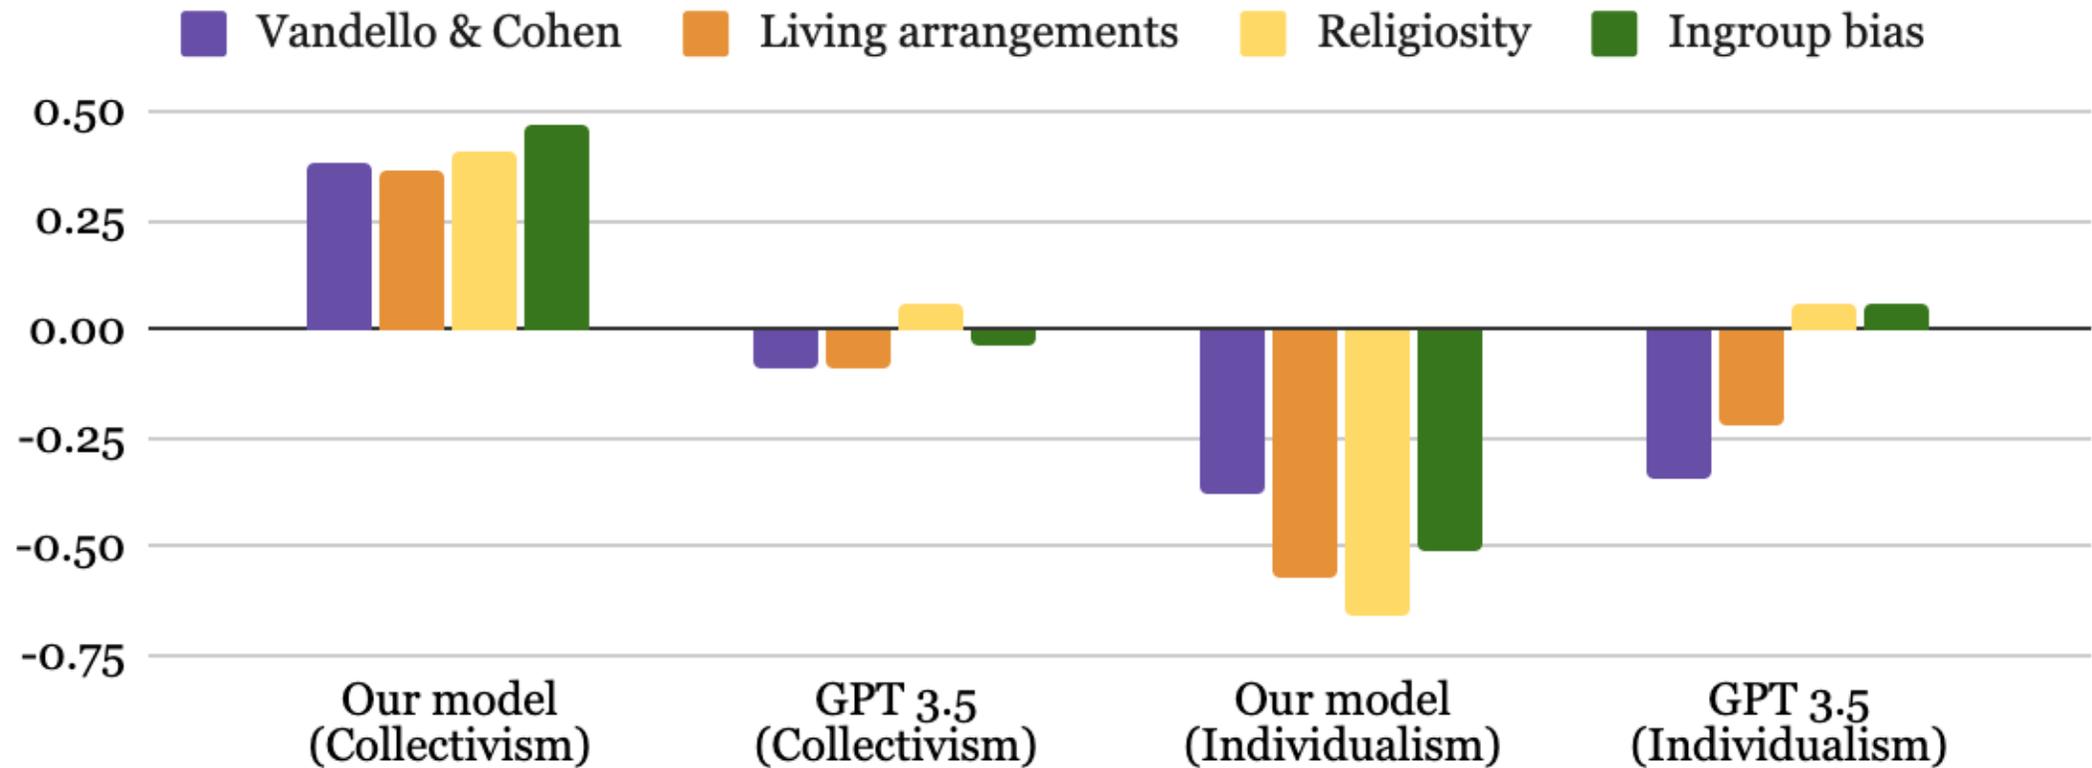
County-level
individualism and
collectivism scores

Individualism (blue) and Collectivism (red) across American counties



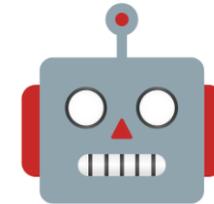
Lexical model vs. zero-shot LLM

Correlation with validated indicators of collectivism

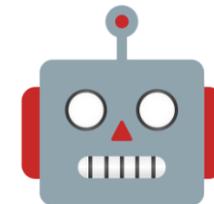


Do LLMs reflect real-world cultural variation?

Nothing beats a slice of New York pizza on a Friday night #NYCeats



Louisiana summers are no joke, y'all. The humidity is on another level.
Stay cool out there, friends! #LouisianaLife



When asked to Tweet given a cultural persona, GPT-3.5 generates ***stereotypical content*** that ***does not match real-word speech patterns.***

Key Takeaways + Future Work

- Lots of work needed to help LLMs understand cultural nuance
- Easily extendible to a multilingual setting
 - Measure culture on a global scale
- Inject understanding of regional variation into LLMs to build culturally adaptable models



With this method, we can (1) measure cultural variation and (2) distill culturally relevant language!

Thank you!
