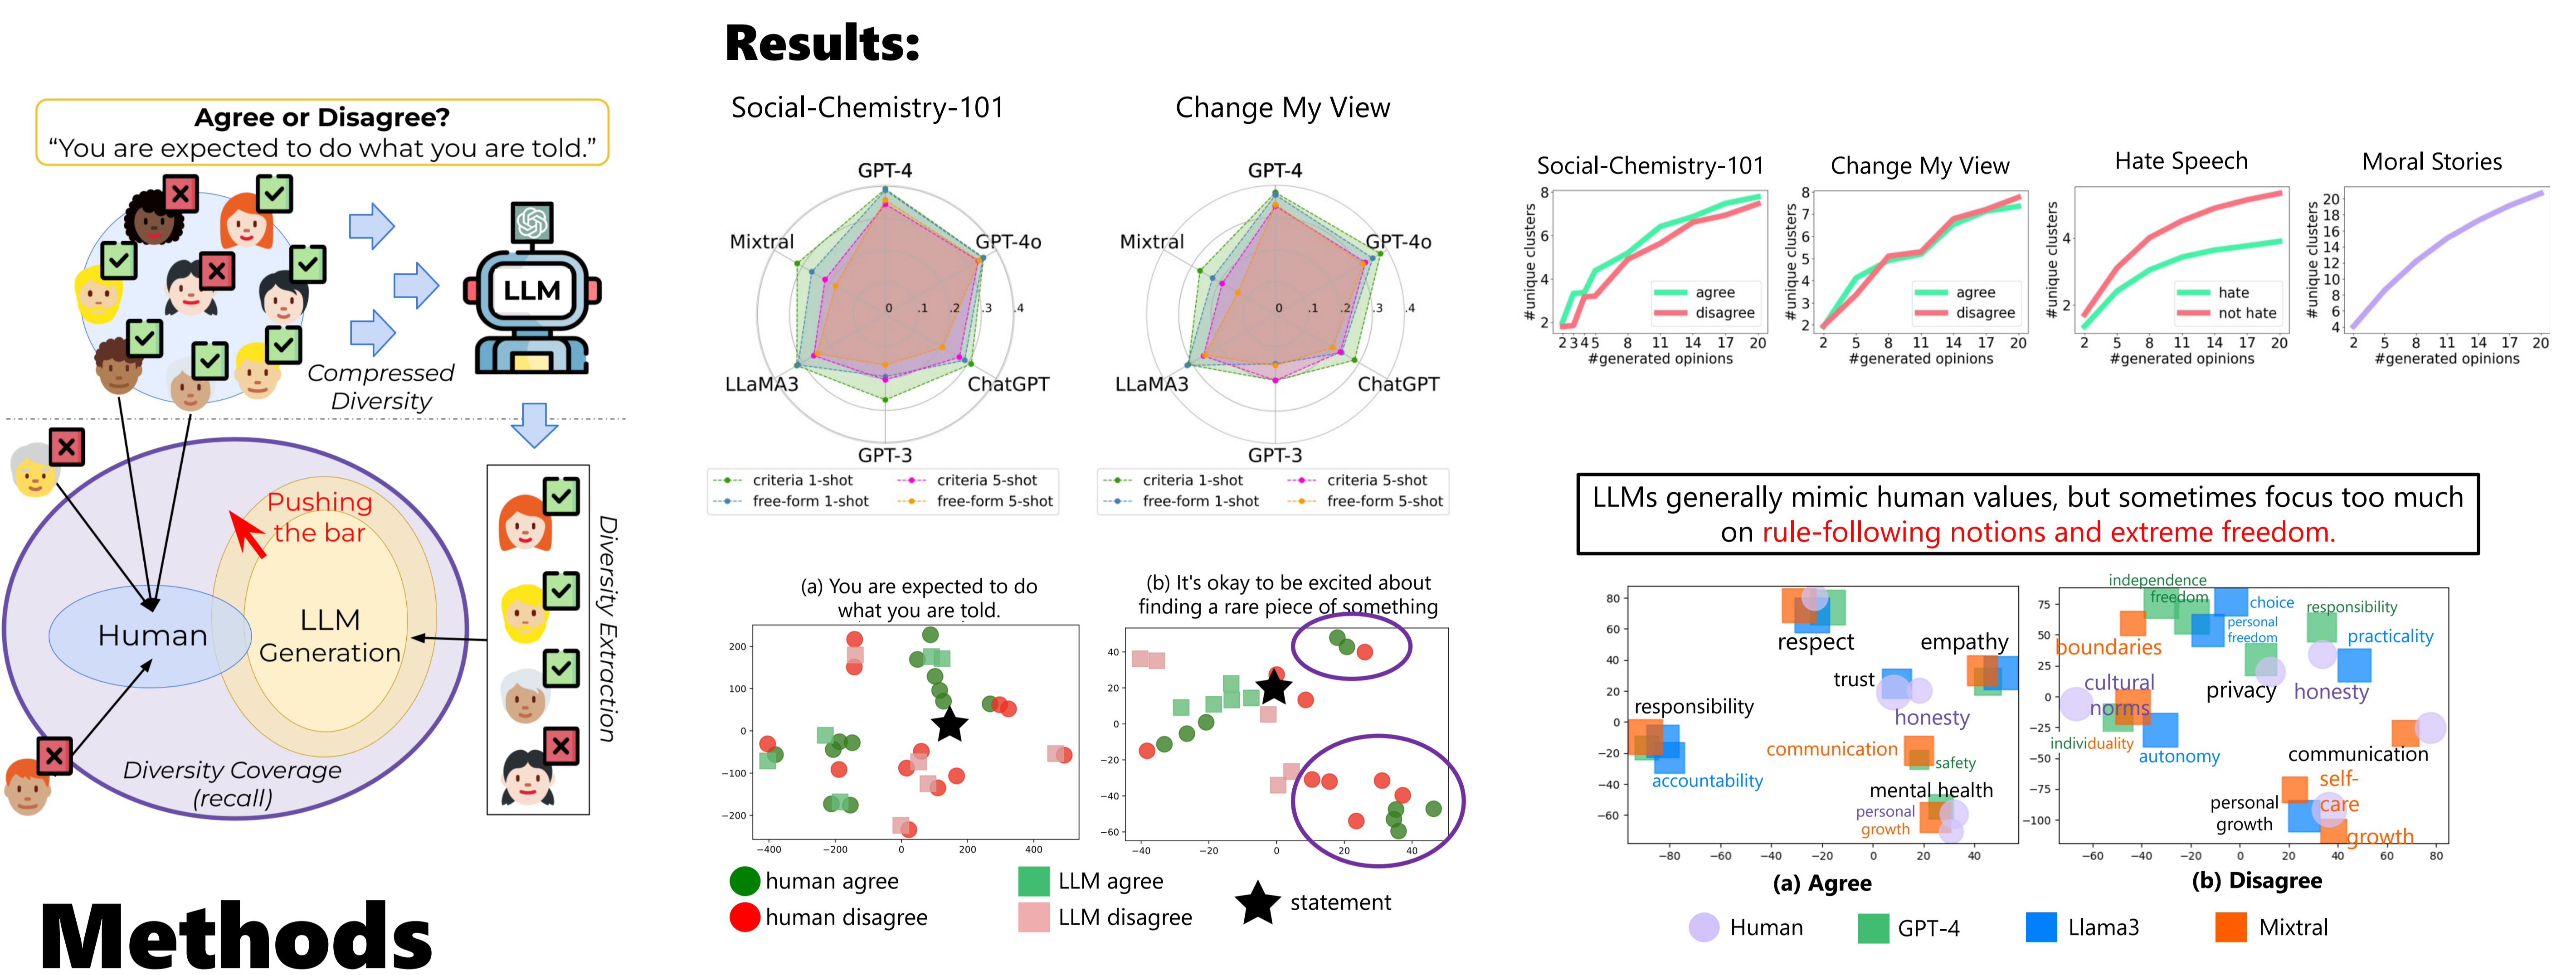


With criteria-based recall prompting, we can measure how much diversity LLMs can extract! 🤖 it's comparable to humans!

How Far Can We Extract **Diverse Perspectives** from Large Language Models?

Background: As collecting diverse human opinions is costly, this leads to a recent trend in human-LLM collaboration for generating diverse data. However, the extent of LLMs' capability to generate maximum diversity remains unexplored.



1 Criteria-Based Prompting

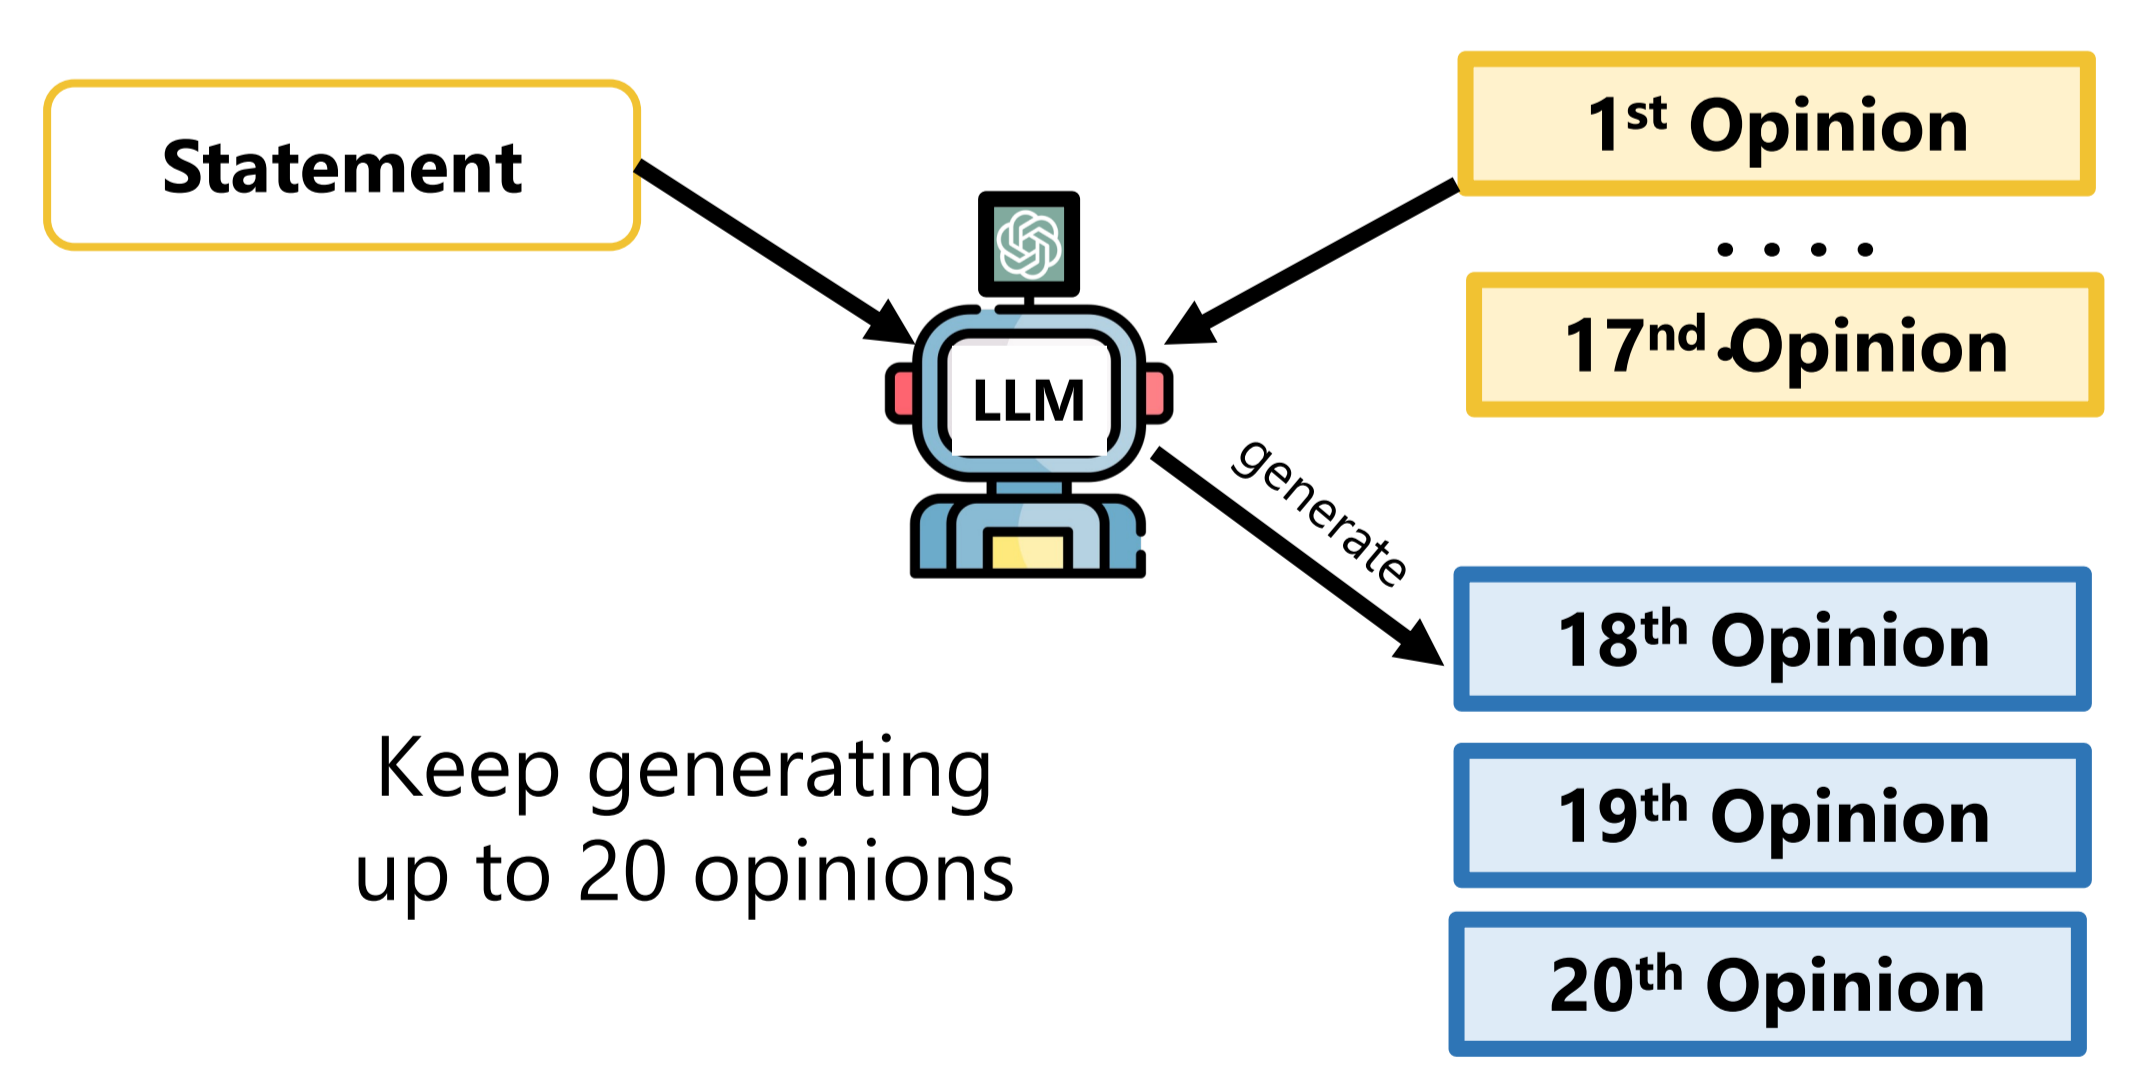
one-shot example

Statement: It's okay to have privacy
Tell me opinions about the statement as many as possible from different people with "Agree" or "Disagree," one-word or one-phrase criteria that is important for their opinions, and explain how they have different opinions
Output:


```
{
  1: {"Stance": "Agree",
      "Criteria": ["personal boundaries", "autonomy"]
      "Reason": "Having privacy allows individuals to establish personal boundaries and maintain their autonomy."}
  2: {"Stance": "...}
  ...
  10: {...}
}
```

Statement: You are expected to do what you are told.

2 Recall Prompting



Conclusion: This work introduces a new task of extracting maximum diversity from Large Language Models (LLMs), and we propose a new prompting method: criteria-based prompting.



Shirley Anugrah Hayati (hayat023@umn.edu)
Minhwa Lee Dheeraj Rajagopal Dongyeop Kang

