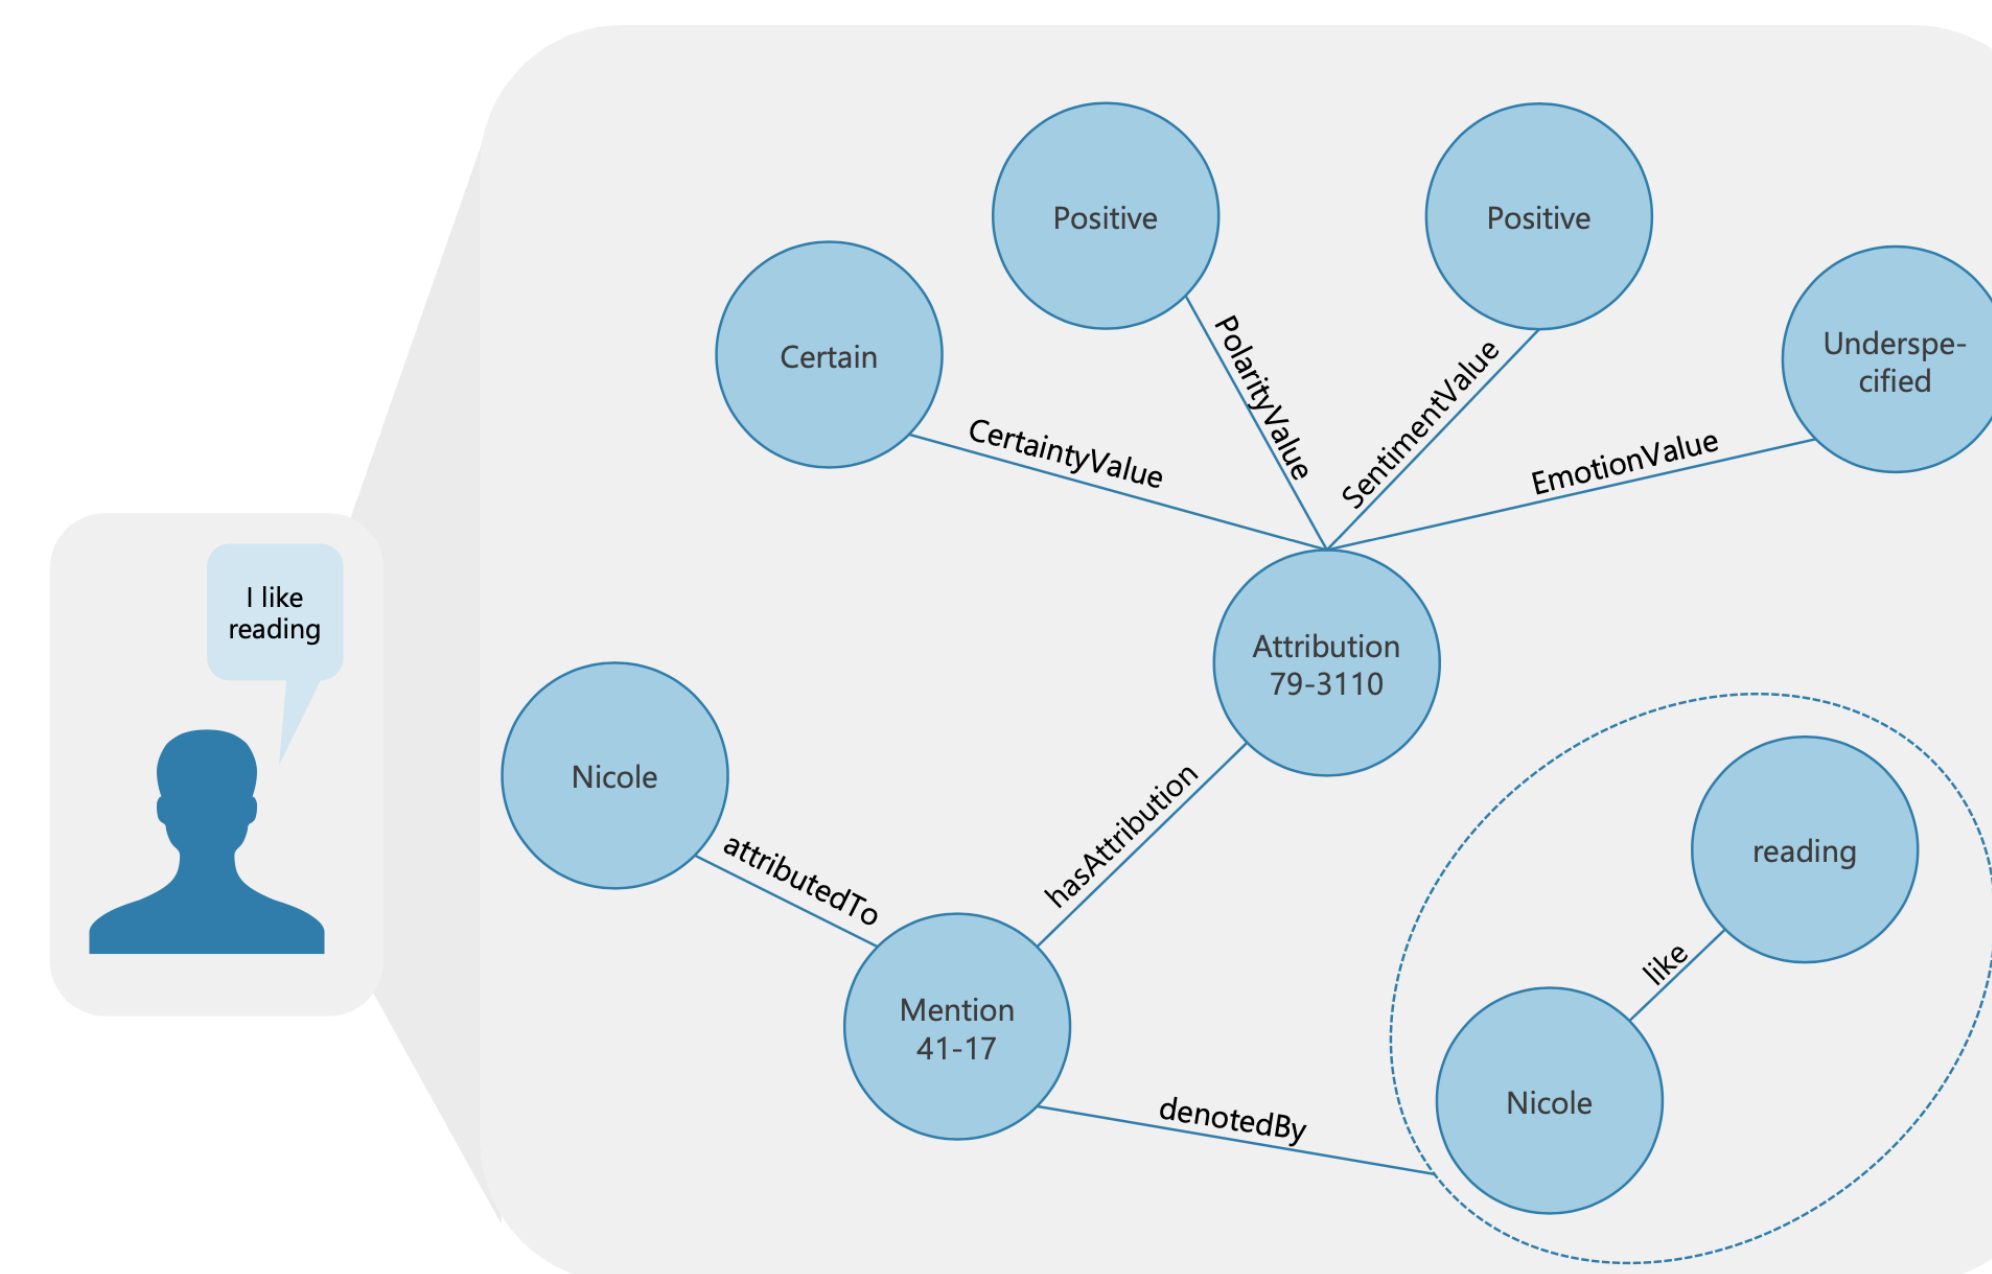


Objectives

- Evaluating the agent's knowledge base considering aspects like source uncertainty, overlap with other sources, conflicting knowledge, timeliness, and the interconnectedness of the knowledge.
- Understanding the user's context, including past dialogues, environmental cues, and user profile.
- Making informed decisions on when and how to communicate knowledge or model uncertainties, setting clear guidelines for when an agent should hold back from responding.

KGs

- Acknowledging the origin or source of a piece of information
- Conversations and people's opinion are dynamic
- → capture shifts in participants opinion or stance over time

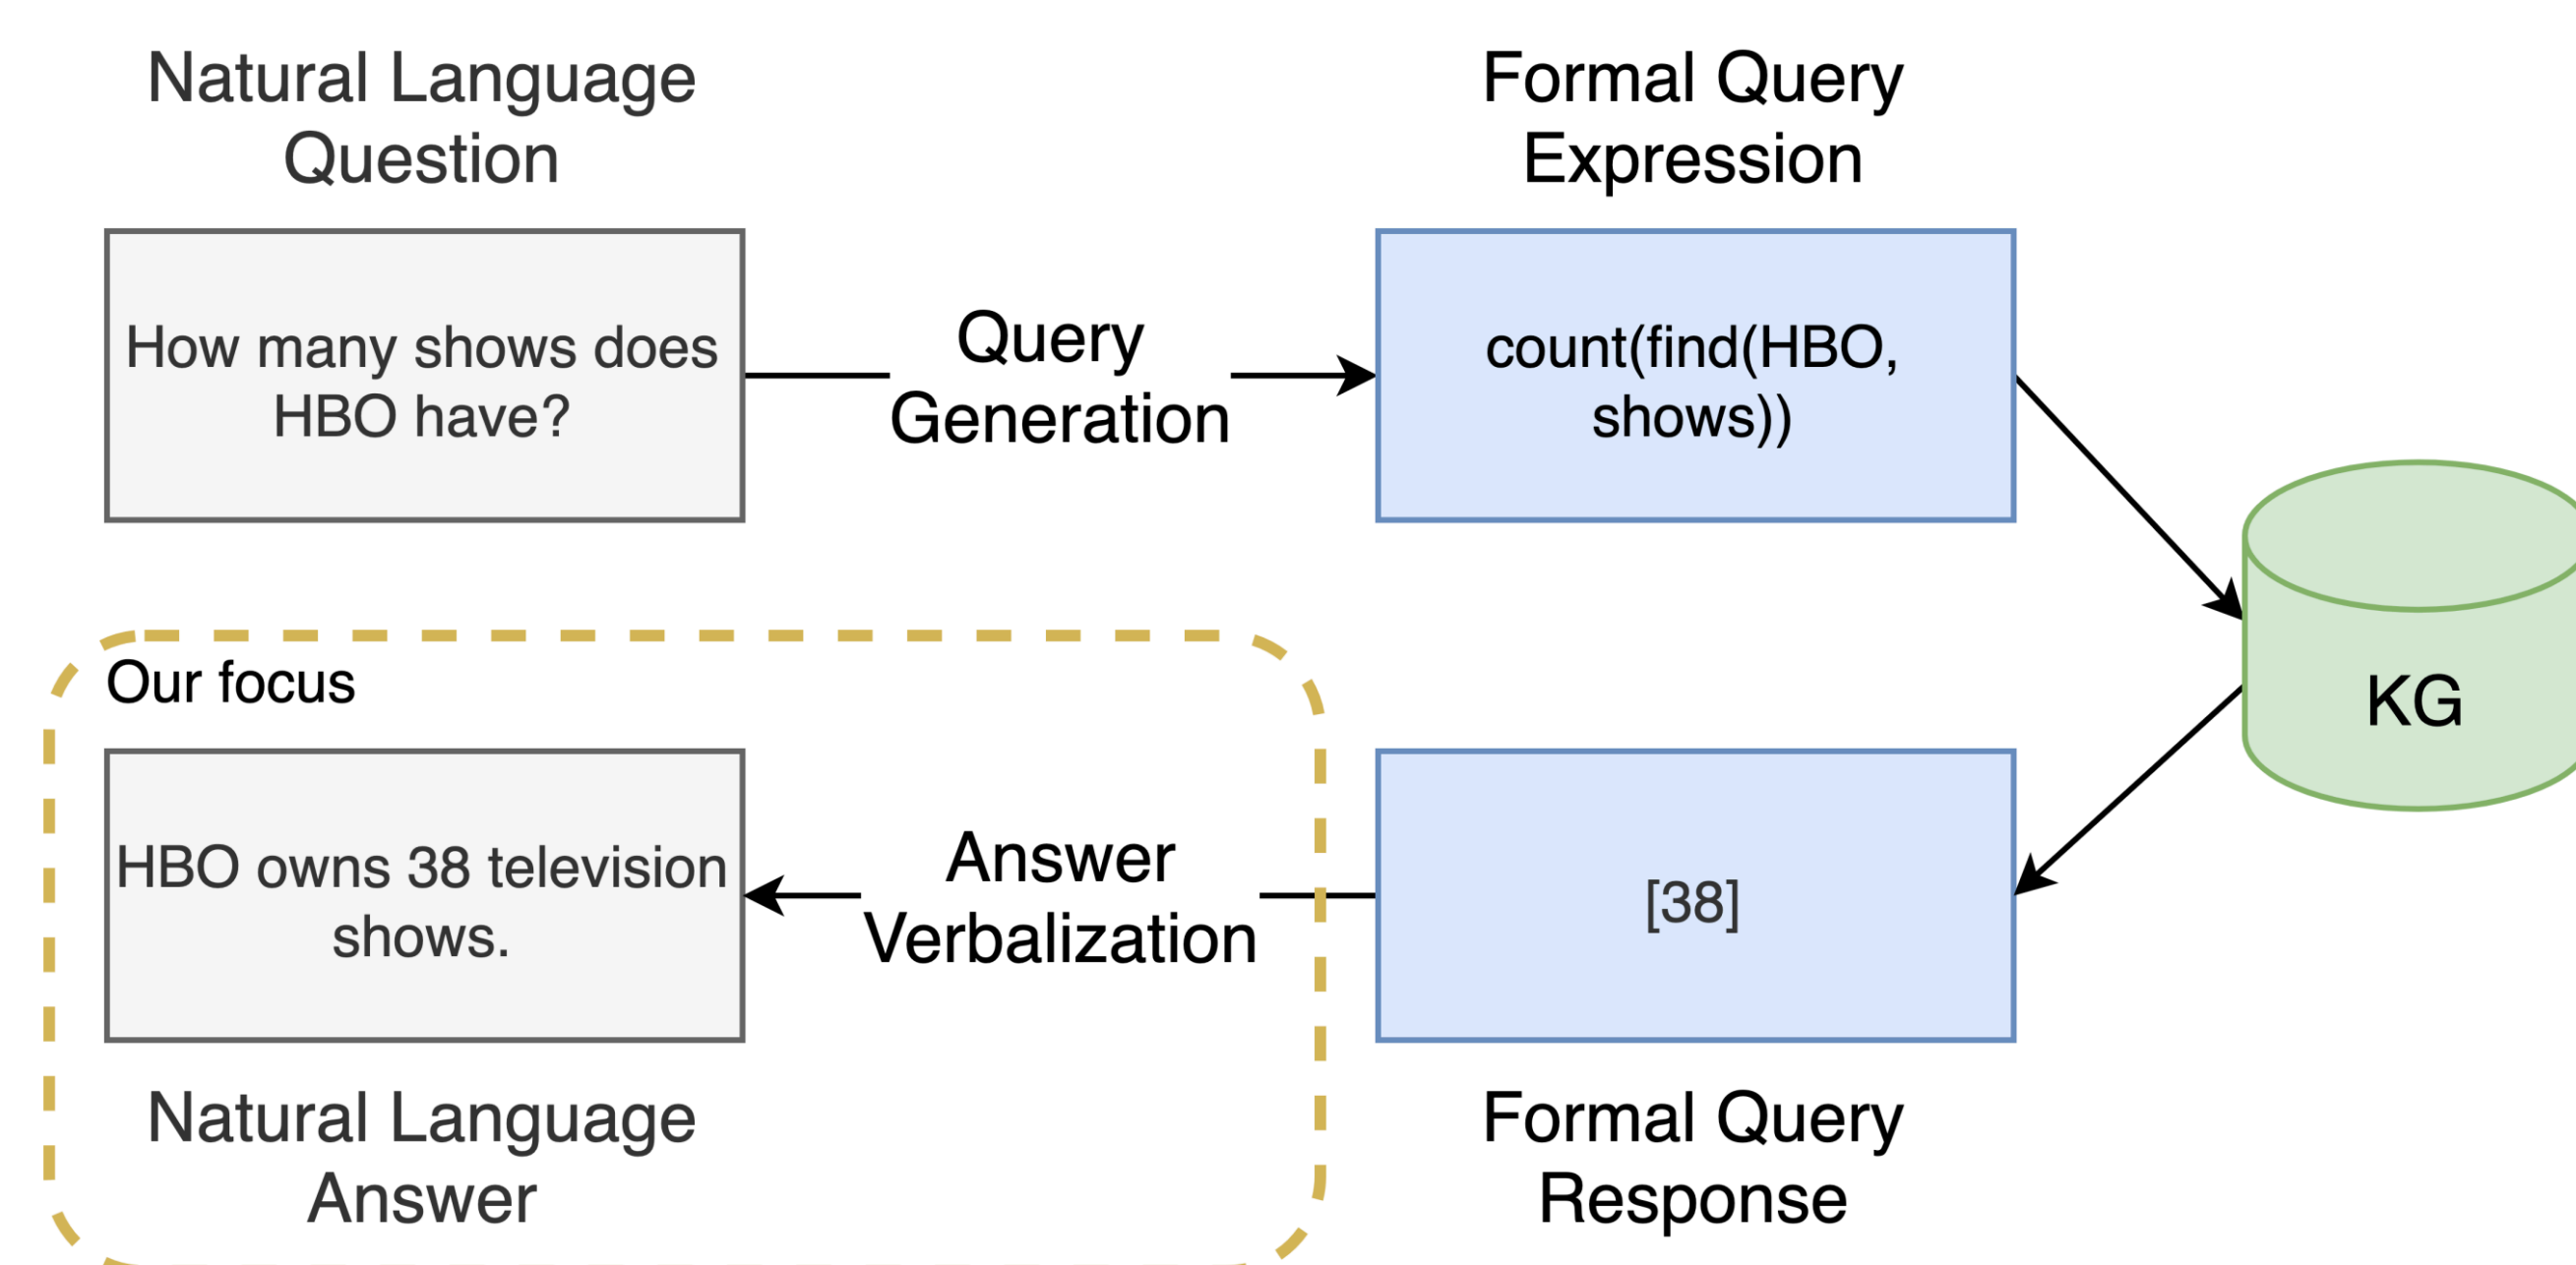


▲ Example of how an utterance is converted into an episodic graph with source perspective values.

Uncertainty

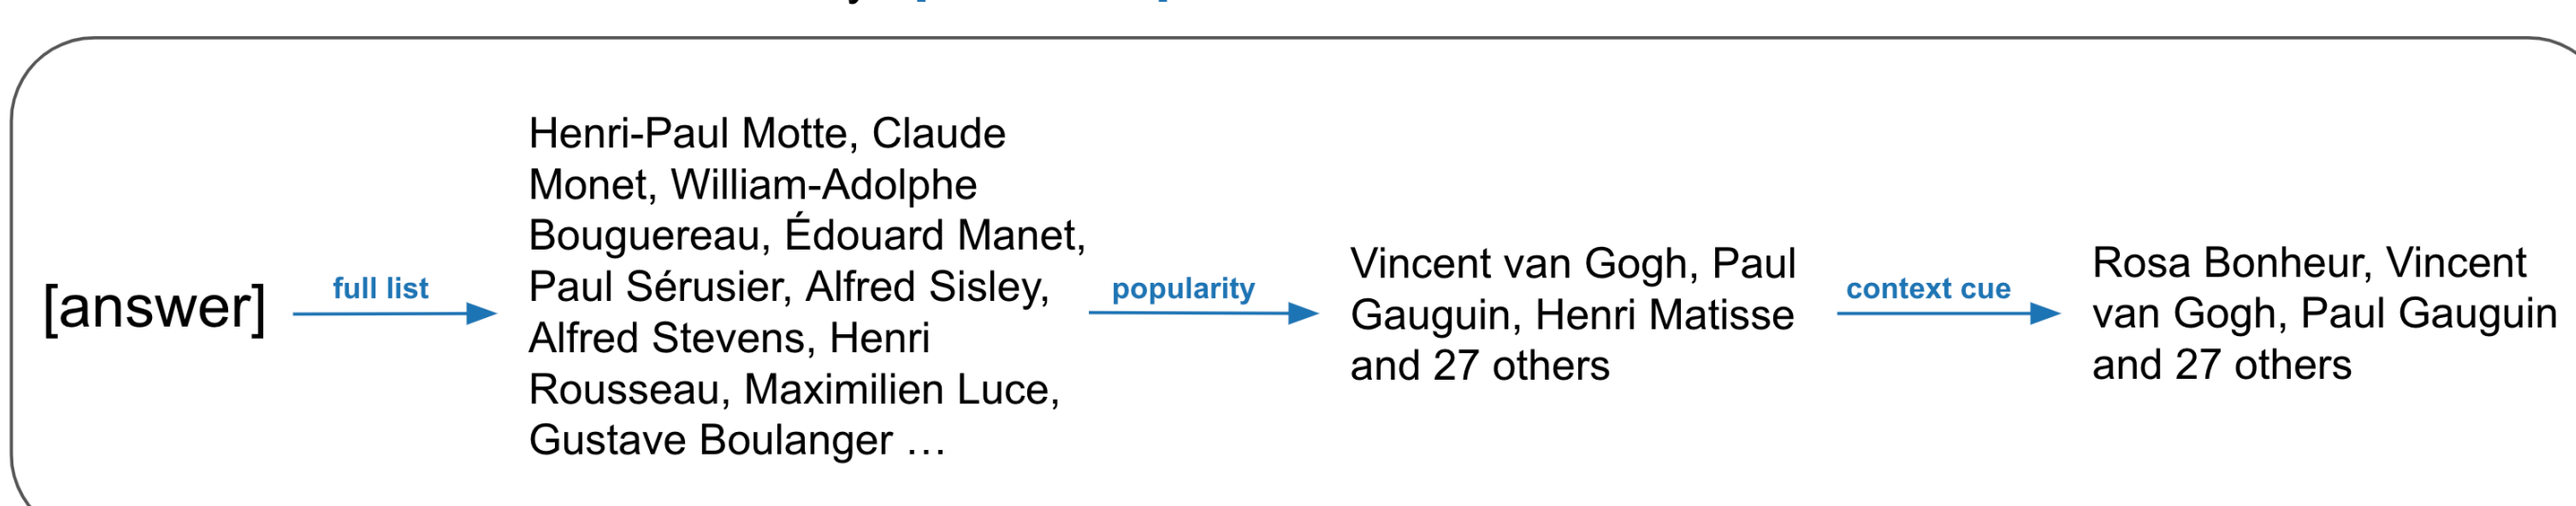
- Uncertainty of the **source**
 - needs to be captured when the utterance is uttered
 - needs to be communicated
- Uncertainty of the **NLG model**
 - needs to be quantified at time of generation
 - needs to be taken into account for answer selection and whether an answer can be given
 - can be communicated, e.g. if a certain threshold is passed

KGQA



- Verbalising and communicating long answers in a human understandable way
- Using Gricean maxims for selection
- Ranking by popularity vs. ranking for contextual factors

Whose work is in the Musée d'Orsay? [female artist]



The artists of the artworks located in the Musée d'Orsay are for example [Rosa Bonheur](#), [Vincent van Gogh](#), [Paul Gauguin](#) and 27 others.

Dialogue

Summarisation

We want more than just a simple summary of factoids / statements.

- **Context:** different people need different summarisations of a dialogue
- **Source attribution:**
 - clearly attribute origin of knowledge / statements
 - capture evolvement of views over time
- **(Un)certainty:**
 - find suitable measurements → calibrate models appropriately
 - when to refrain from using a prediction
 - communicate it in the summary

Future

- KG-Dialogue aligned dataset
- Local/culture-specific knowledge in LLMs
- Multilingual certainty calibration
- Mapping uncertainty measures to types of hallucinations

References

1. Krause, L., & Vossen, P. (2020). When to explain: Identifying explanation triggers in human-agent interaction. Proceedings of the 2nd Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence.
2. Krause, L., Sommerauer, P., & Vossen, P. (2022). Towards More Informative List Verbalisations. Proceedings of the 1st Workshop on Knowledge Graph Summarization.
3. Krause, L., Tufa, W. T., Báez Santamaría, S., Daza, A., Khurana, U., & Vossen, P. (2023). Confidently Wrong: Exploring the Calibration and Expression of (Un)Certainty of Large Language Models in a Multilingual Setting. Proceedings of the 1st Workshop on Multimodal, Multilingual NLG (MM-NLG).