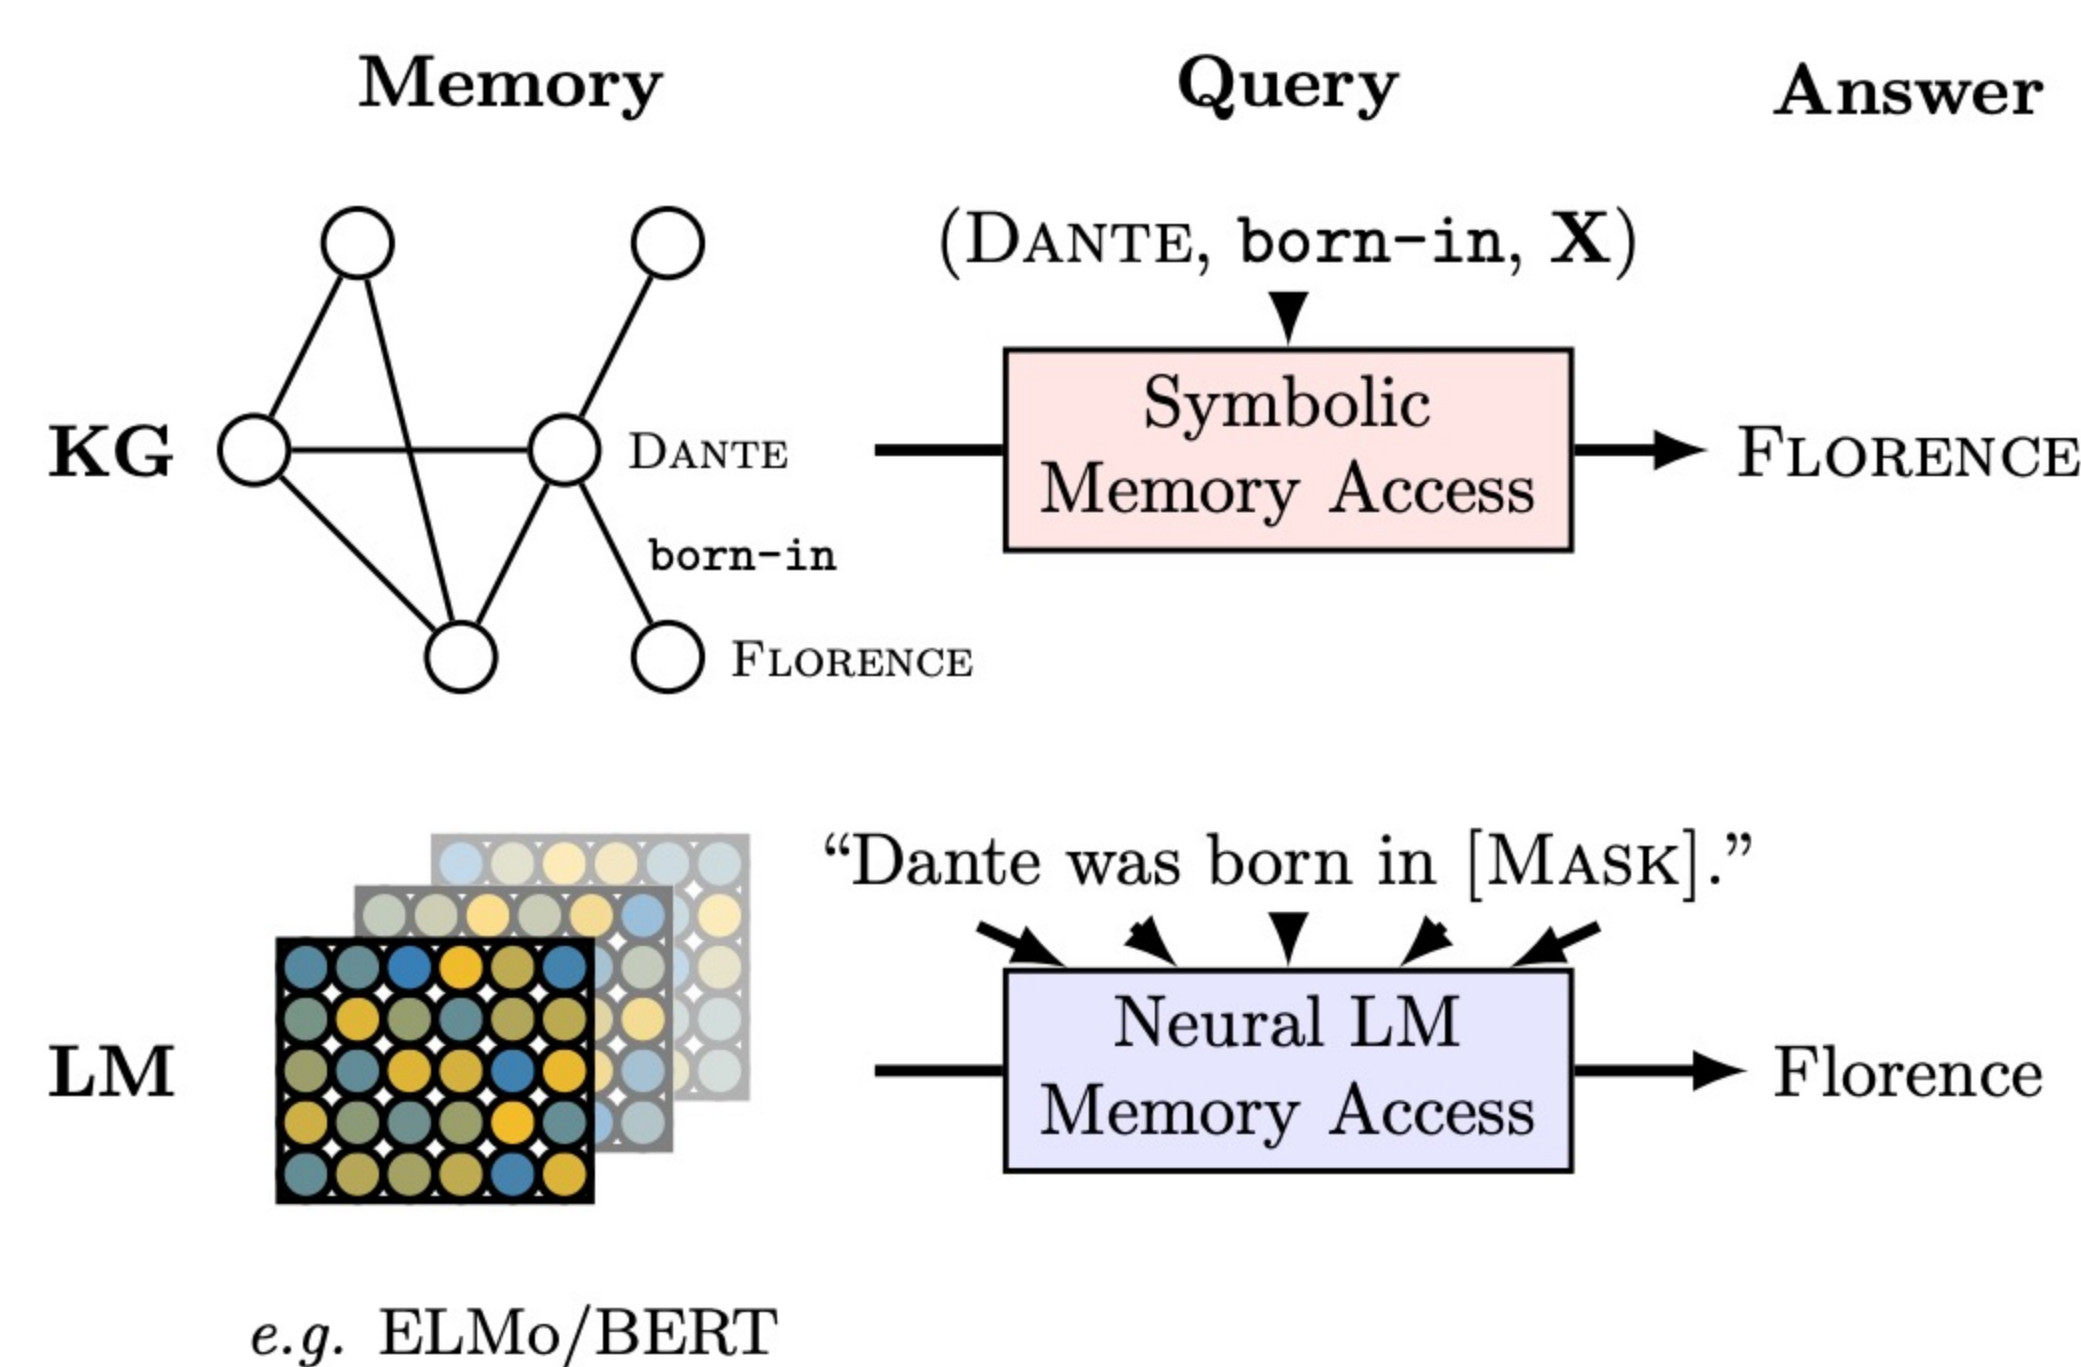


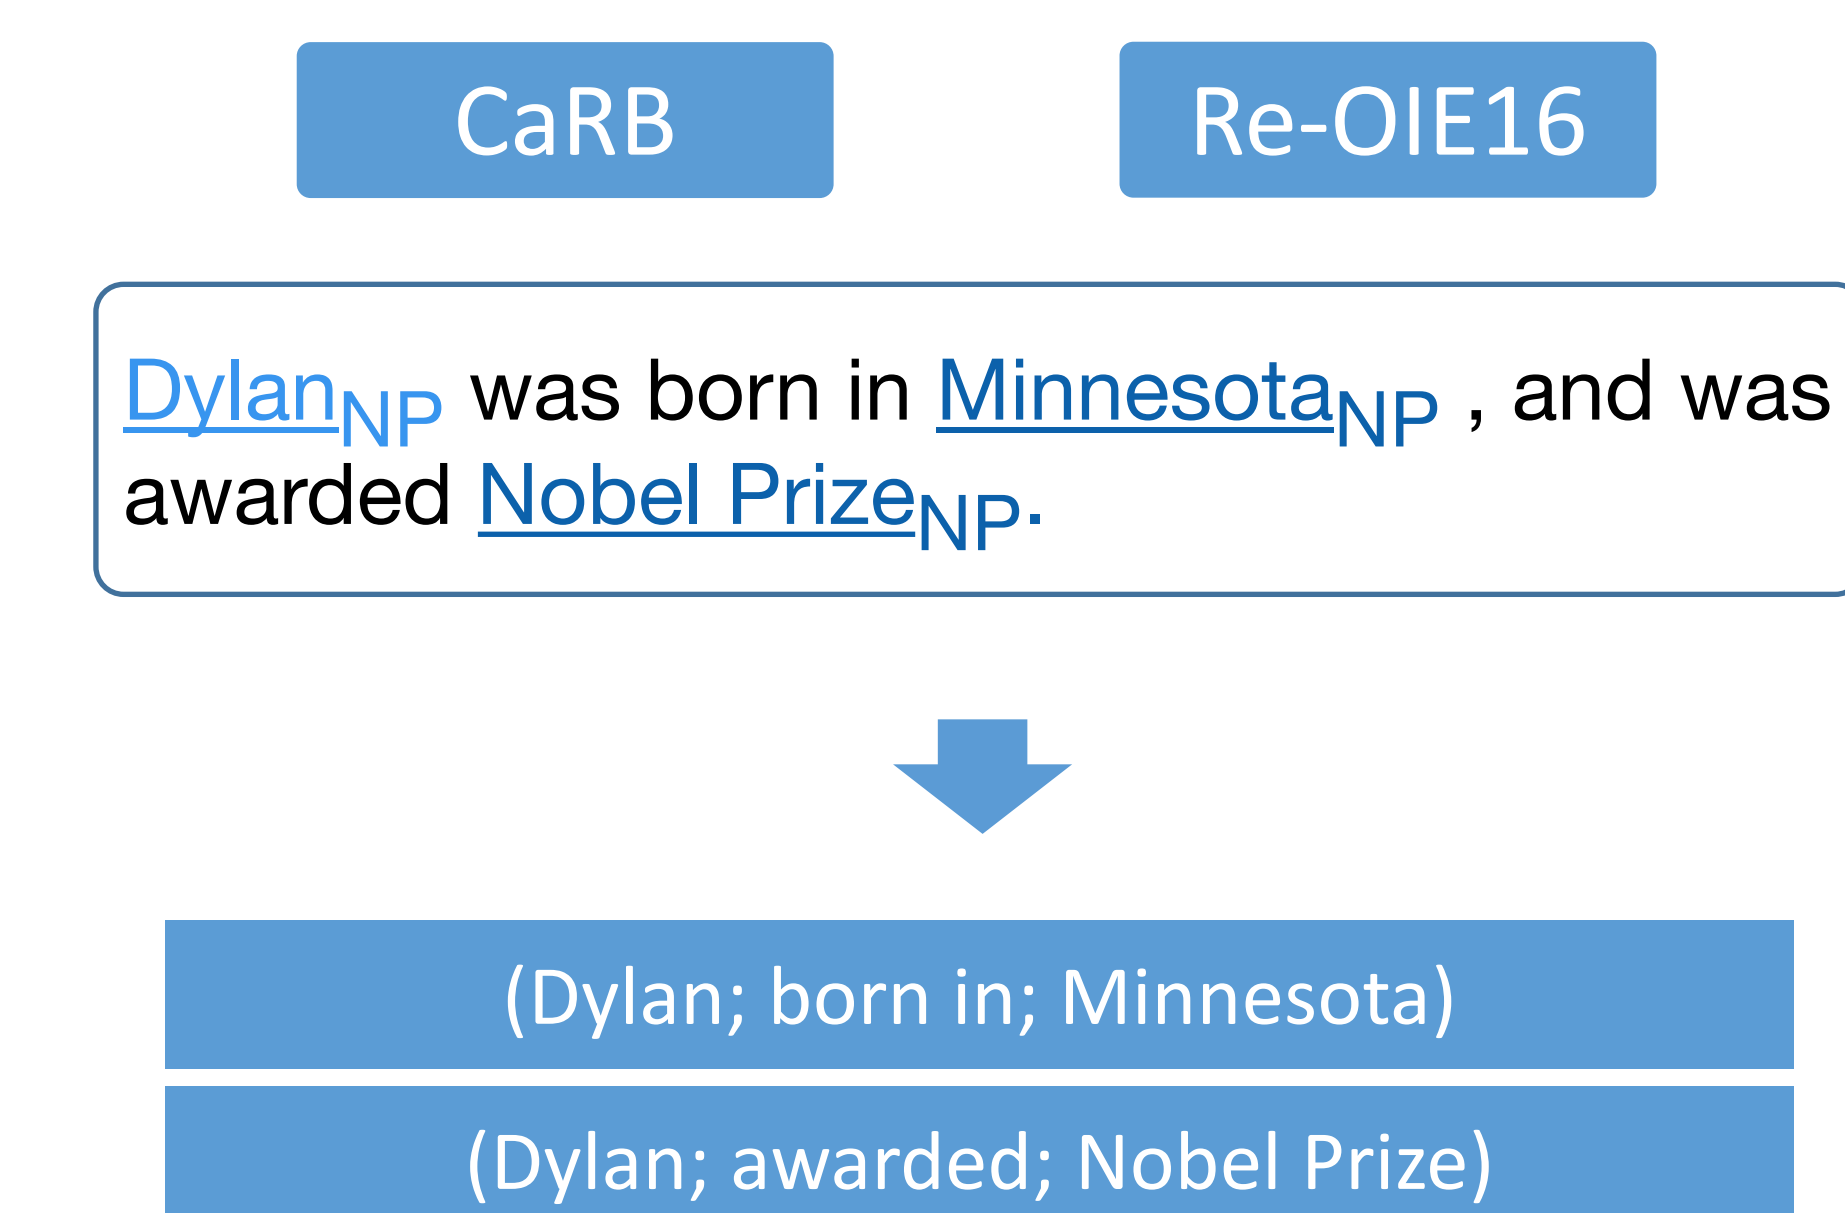
## Our IELM Benchmark

### LAMA – LMs capture Fixed Relational Knowledge




- **Data Source:** Wikidata
- **Construction:** Automatic
- **Evaluation:** Exact Match
- **Statistics:** 34K triples, 41 relations types

### Standard OIE - Open Relational Information



- **Data Source:** News corpus
- **Construction:** Manual Annotation
- **Evaluation:** Lexical Match
- **Statistics:** 4.2K triples, 4.2k predicates

### Factual OIE (New) - Open Relational Knowledge




- **Data Source:** News corpus & Wikipedia
- **Construction:** Manual & Distant Supervision
- **Evaluation:** Exact Match
- **Statistics:** 27M triples, 1,197 relation types

Dataset	#Triples	#Args	#Preds	#Doc
CaRB	1,508	3,328	1,506	595
Re-OIE16	2,715	6,226	2,715	641
TAC KBP-OIE	27,655	39,661	41	3,877,207
Wikidata-OIE	27,368,562	6,047,494	1,156	6,047,494


## Results

- We test the OIE performance of 6 pre-trained LMs (BERT and GPT-2 families) and 14 OIE systems on IELM benchmark.

### Standard OIE CaRB



### Factual OIE Wikidata



- **Observations:**
  - Language models have captured open relational knowledge during their pretraining.
  - Zero-shot language models on factual OIE is better than supervised methods.
  - The larger the LMs, the better zero-shot OIE performance they have.



Get our code

We enable zero-shot pretrained LMs for OIE by encoding the arguments in the input and decoding predicates using the parameters (attention) of pre-trained LMs.

