


Entity Cloze By Date: What LMs Know About Unseen Entities

Yasumasa Onoe, Michael J.Q. Zhang, Eunsol Choi, Greg Durrett

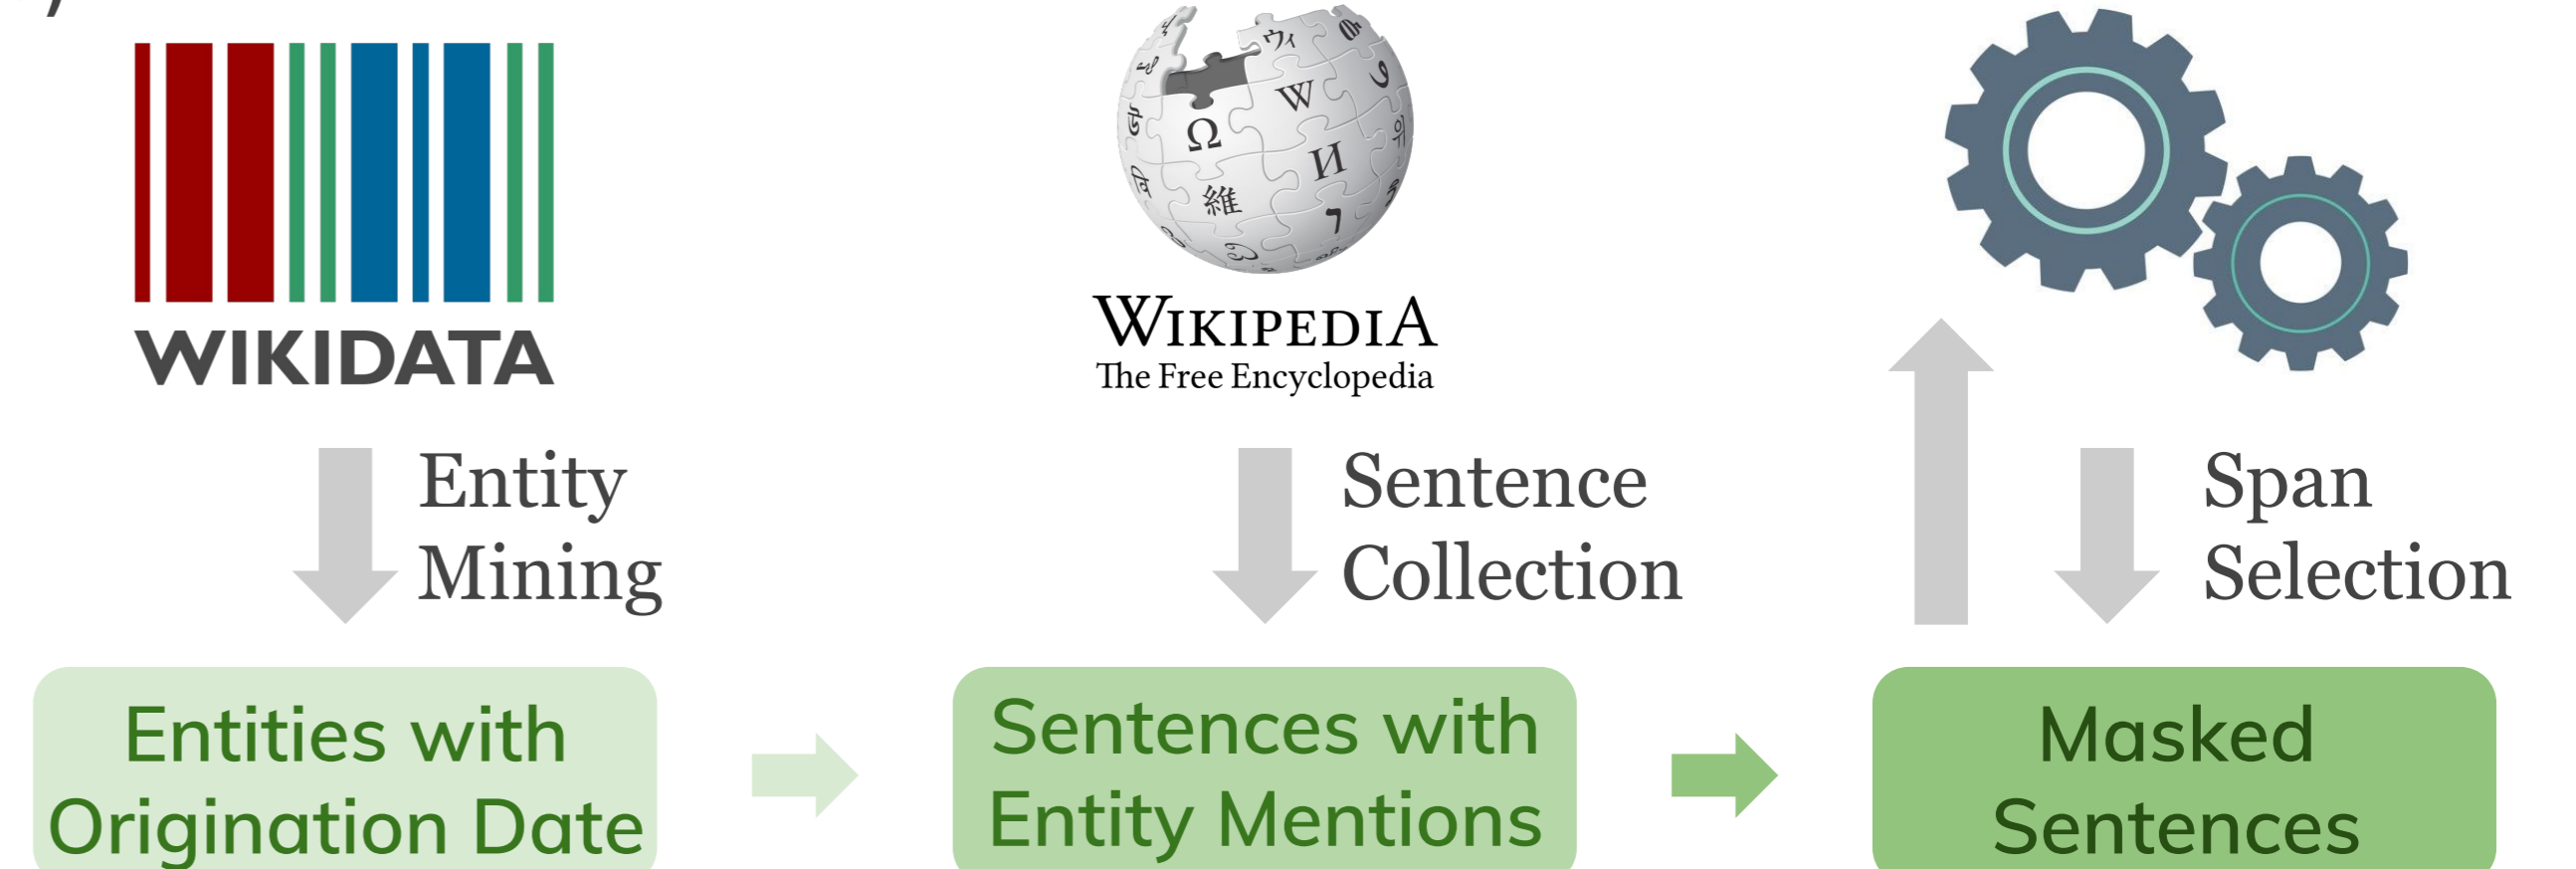
Motivation

- LMs have rich entity knowledge about their training data, but as new entities emerge, this knowledge becomes outdated. How can we test what models know about new entities?
- Existing benchmarks usually test KB triples. How can we test a broader set of inferences about entities?



Data Collection

- Date Indexing**
ECBD is a cloze-style task. Each **masked span** is associated with an **entity** which is paired with its **origination year**.
- Diverse Masked Spans**
NP spans primarily represent relational knowledge. Random spans cover other types of knowledge (e.g., actions an entity can take).



What do large LMs know about new entities that appeared after pre-training?

The **[Windows 11]** name was accidentally released in **[MASK]** in June 2021.
Entity from 2021 NP: an official Microsoft support document

The **[GPT-2]** model has **[MASK]**, and was trained on 8 million web pages.
Entity from 2019 NP: 1.5 billion parameters

[Austin FC] are the operators of a newly-**[MASK]**.
Entity from 2018 Random: built stadium at McKalla Place

[Hurricane Irma] affected multiple states in the **[MASK]**.
Entity from 2017 Random: South, especially Florida

Task

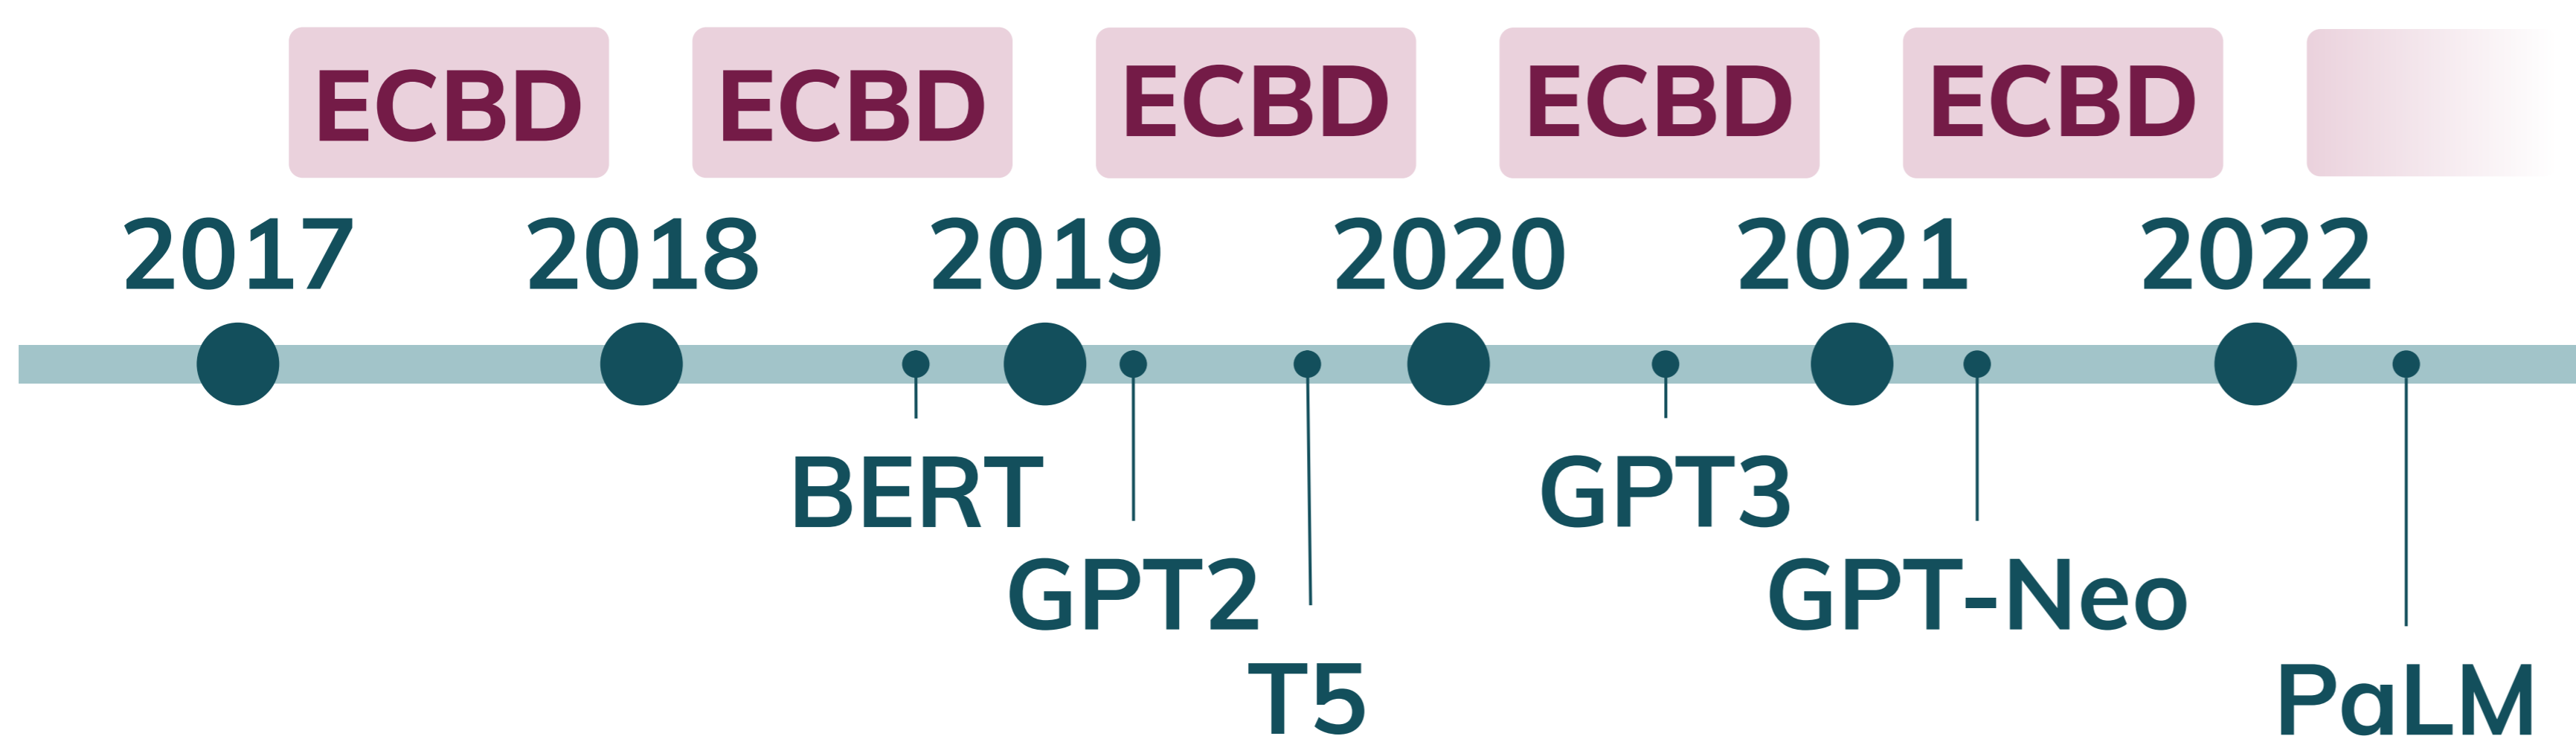
Given a cloze sentence with a **new entity**, predict **masked tokens** (measure the model's perplexity)

Seen/Unseen Entities

Test examples are grouped by the origination year. It's easy to test LMs on **OLD/NEW** entities.

Automatically constructed; we can update ECBD easily!

Can test if knowledge editing approaches successfully inject information




Data Statistics

Size

	# Sent.	# Ent.	Span Len.	Span V.
ECBD	35k	2,106	2.9	19,542
Popular*	8k	1,910	2.9	8,644

* Set of existing entities to use a baseline.

New Entity Types



Experiments

What do LMs know about seen/unseen entities?

- Compare T5-Large with different inputs.
 - No Ent: remove entity mention, Definition: add definition from EN Wikipedia
- Measure perplexity on the masked spans

Observation

- Newer entities may be harder.
- Adding definition is particularly helpful for new entities – **Entity knowledge matters for this task!**

