

Retrieval / Knowledge-Intensive NLP









1









\wedge	What	do	the	HMs	"know?	
	VVIIdL	uυ	une	LLIVIS	KNOWS	

- iPod Touch is produced by ____.
- London Jazz Festival is located in _____.
- Dani Alves plays with _____.
- Carl III used to communicate in _____
- Ravens can _____.

What do the LLMs "know?

- iPod Touch is produced by Apple.
- London Jazz Festival is located in London.
- Dani Alves plays with Santos.
- Carl III used to communicate in German.
- Ravens can fly.

(Petroni et al., 2019)

What do the LLMs "know?

Lots of knowledge from language modeling

Issues:

- Coverage: was the fact in the training set?
- Frequency of facts: has the LM "seen" it enough times?Prompt sensitivity: if we reword it, will the answer change?























Continuous Cache	# Retrieval tokens	Granularity	Retriever training Frozen (LSTM)	Retrieval integration Add to probs
kNN-LM	O(109) Type	1 Token	Frozen (Transformer)	Add to probs
SPALM	O (10 ⁹)	Token	Frozen (Transformer)	Gated logits
DPR	O (10 ⁹)	Prompt	Contrastive proxy	Extractive QA
REALM	O (10 ⁹)	Prompt	End-to-End	Prepend to prompt
RAG	O (10 ⁹) Type :	Prompt	Fine-tuned DPR	Cross-attention
FID	$O(10^{9})$	Prompt	Frozen Dpr	Cross-attention
Emdr ²	$O(10^{9})$	Prompt	End-to-End (EM)	Cross-attention
RETRO (ours)	O (10 ¹²)	Chunk	Frozen (BERT)	Chunked cross-attention
Type 1: Te	oken-level Retrieva	l (mainly) fo	r <i>LM</i> – augmenting p	rediction of next token
Type 2: Pa:	ssage-level Retrieval (mainly) for Q2	1 – retrieving passages r	elevant to the question



\bigwedge	Prompt-Level Granularity
	$\label{eq:response} \begin{array}{ c c } \hline \mbox{Traing} \\ \hline \mbox{Traing} \\$
	Dense Passage Retriever (Karpukhin et al., 2020)





RETRO	
INTE The Dure film was released in	

	RETRO					
	Key (BERT sentence embedding)	Value (text. neighbor and completion chunks. Each up to 64 tokens in length)				
Database		Dune is a 2021 American epic science fiction film directed by Denis Villeneuve	NEIGHBOR			
		It is the first of a planned two-part adaptation of the 1985 novel by Frank Herbert	COMPLETION			
		Dune is a 1965 science fiction novel by American author Frank Herbert	NEIGHBOR			
~~		originally published as two separate serials in Analog magazine	COMPLETION			















