


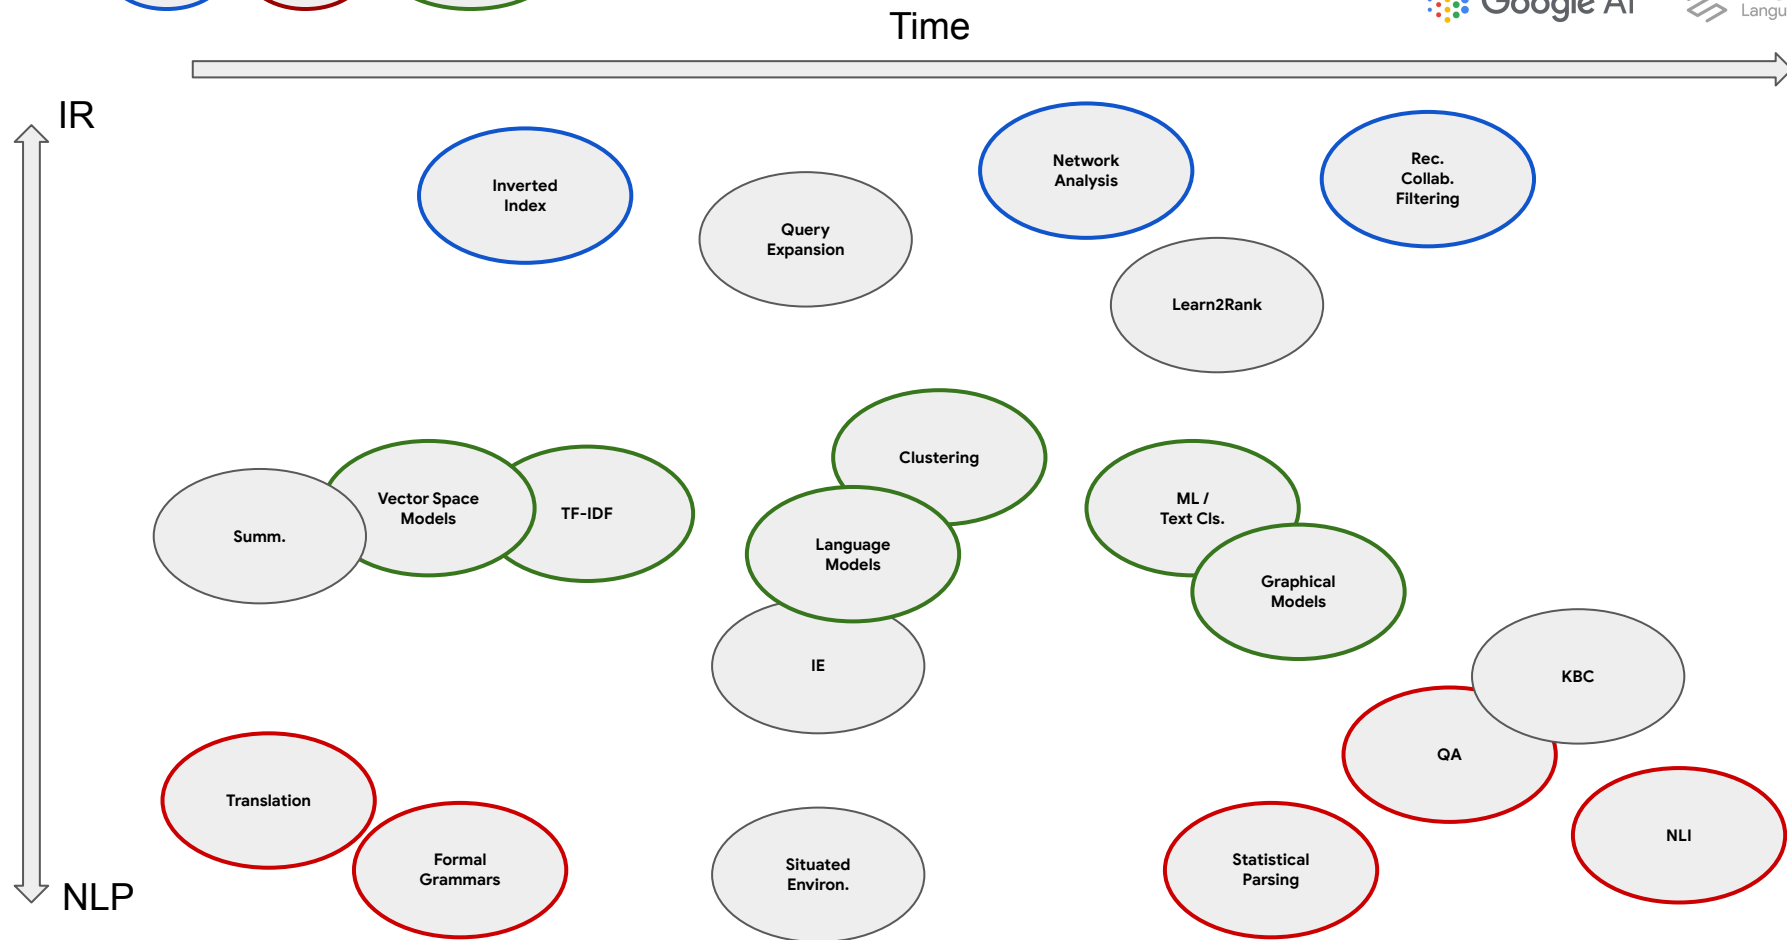
# NLP & IR

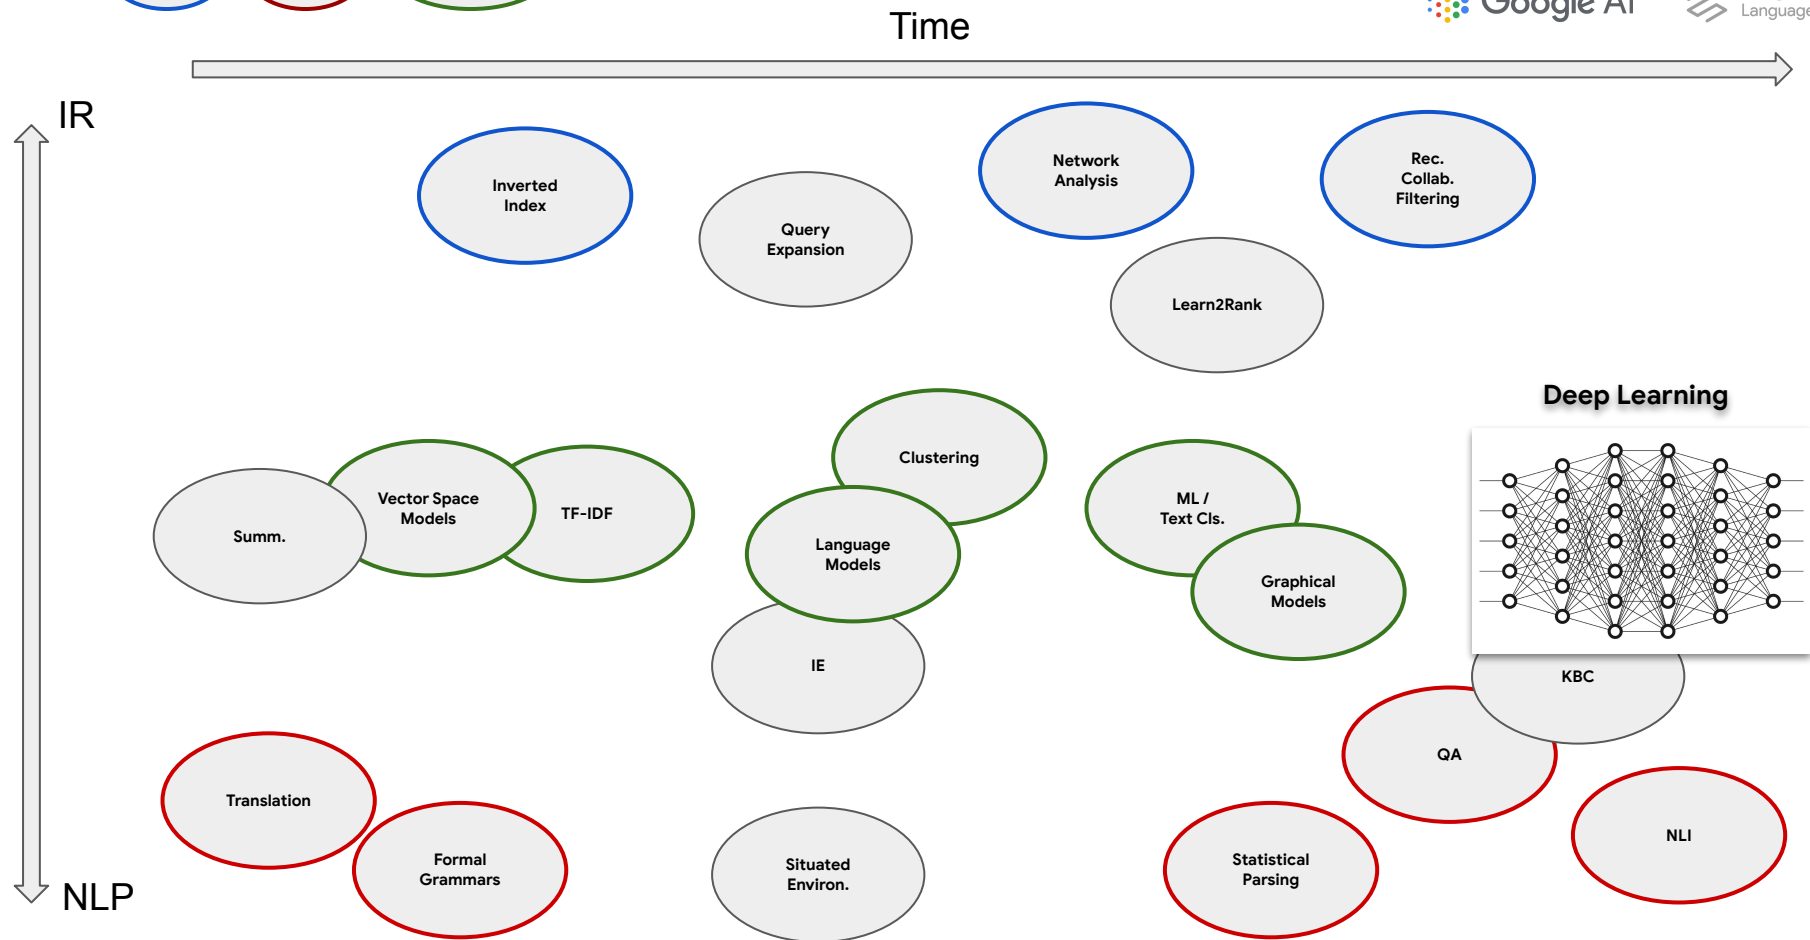
## How Deep Learning has Bridged the Gap

Ryan McDonald

*Material from a number of Google research projects*







**$f(\text{text}, \text{text})$**

**$f(\text{text}, \text{entity})$**

**$f(\text{text}, \text{image})$**

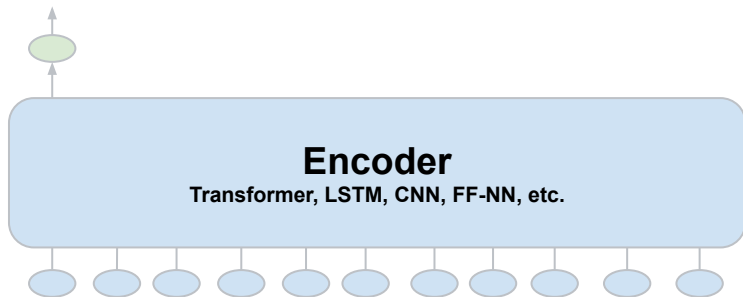
**...**

**$f(\text{text}, \text{object})$**




# $f(\text{text}, \text{text})$

Entail



[CLS] A boy and his mother and father are at the beach [SEP] A family is doing something outside

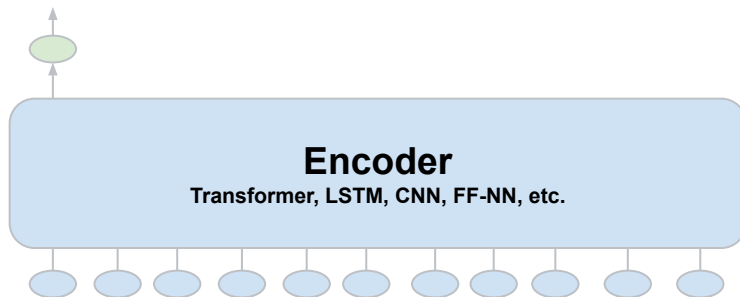
Entail



A boy and his mother and father are at the beach


A family is doing something outside

Contradict



[CLS] A man inspects the uniform of a figure in some East Asian country. [SEP] The man is sleeping

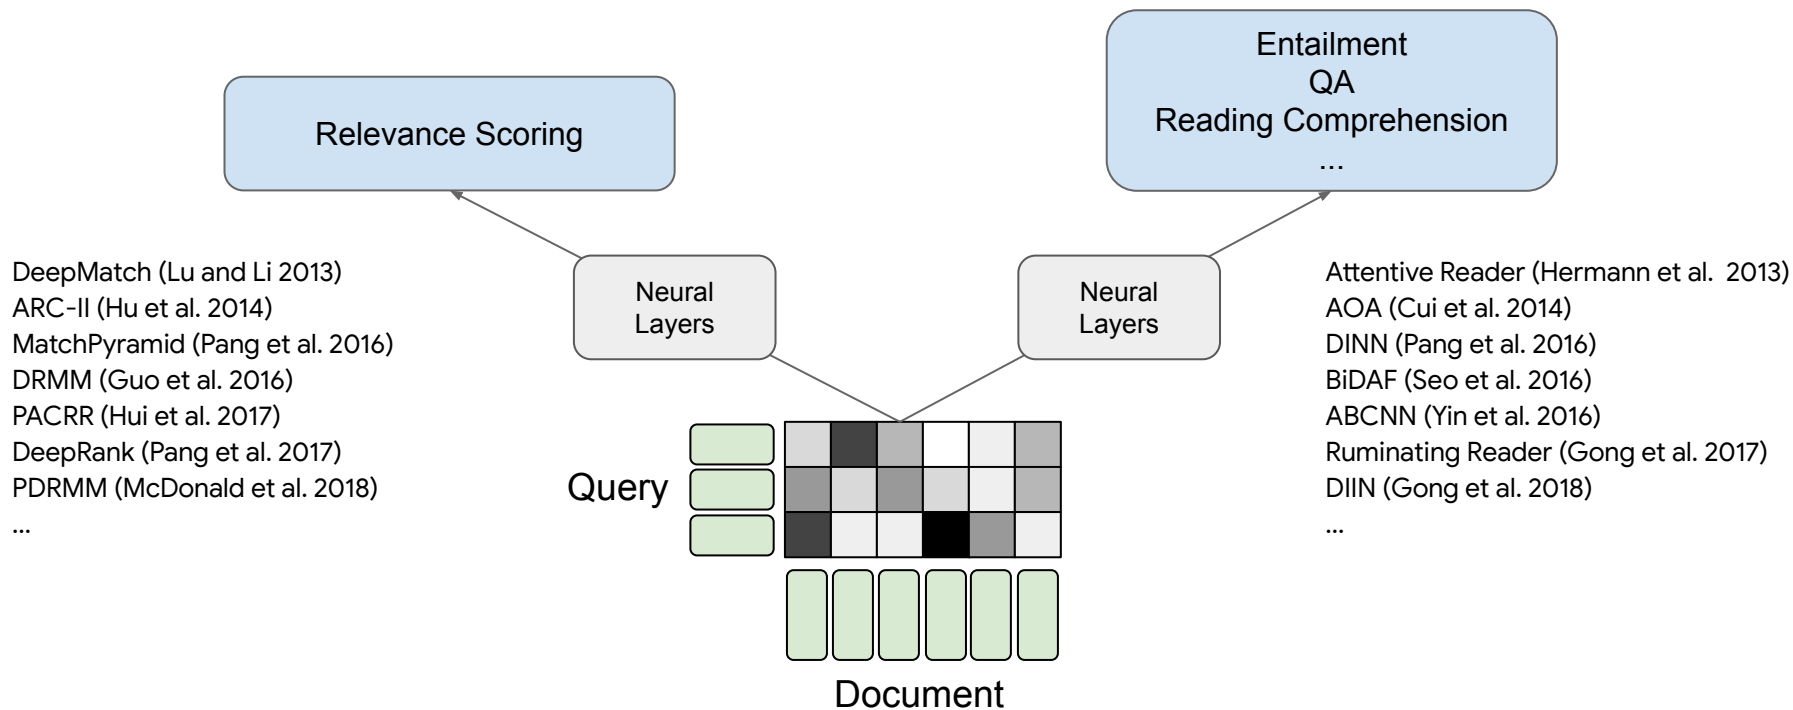
Contradict



A man inspects the uniform of a figure in some East Asian country.

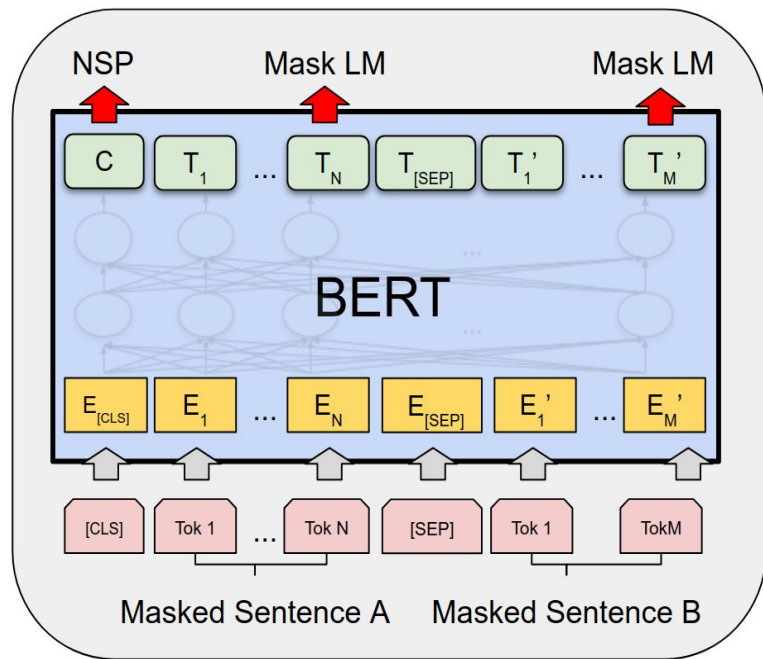
The man is sleeping

# Cross-attention (AKA one-tower; AKA interaction-based)



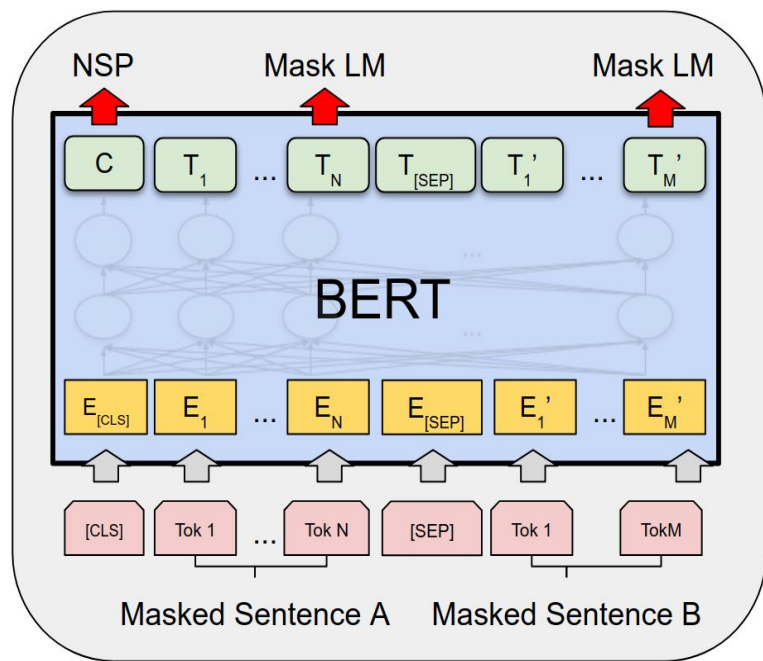
# BERT: Transformers + Pre-training + Fine-Tuning

## Pre-Training

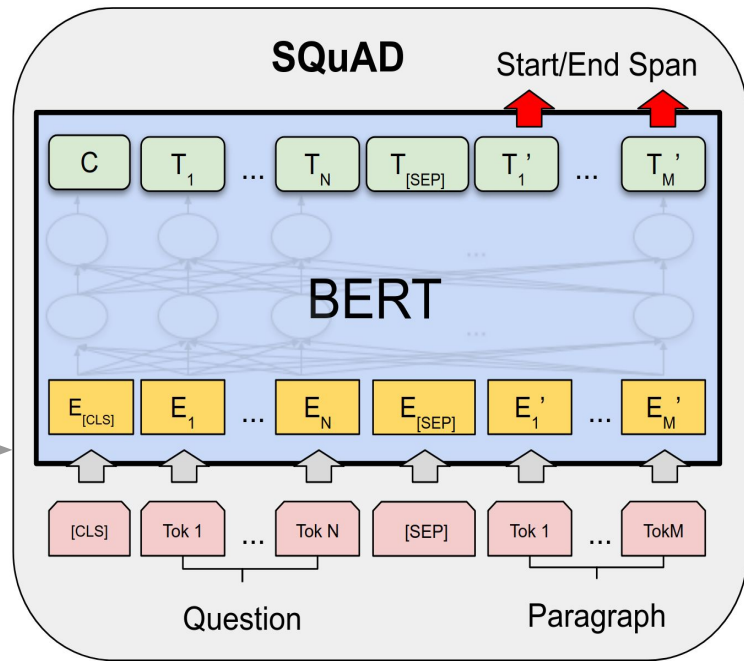


# BERT: Transformers + Pre-training + Fine-Tuning

## Pre-Training



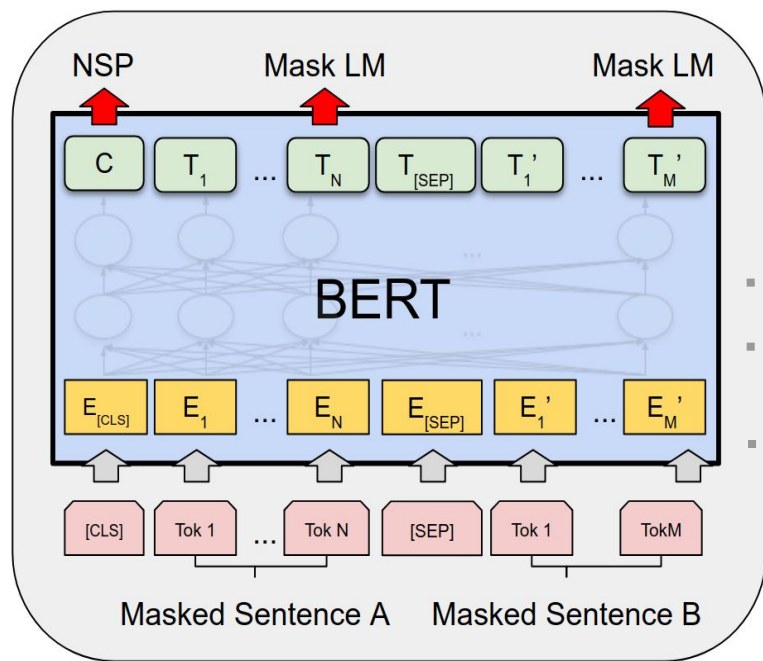
## Fine-Tuning



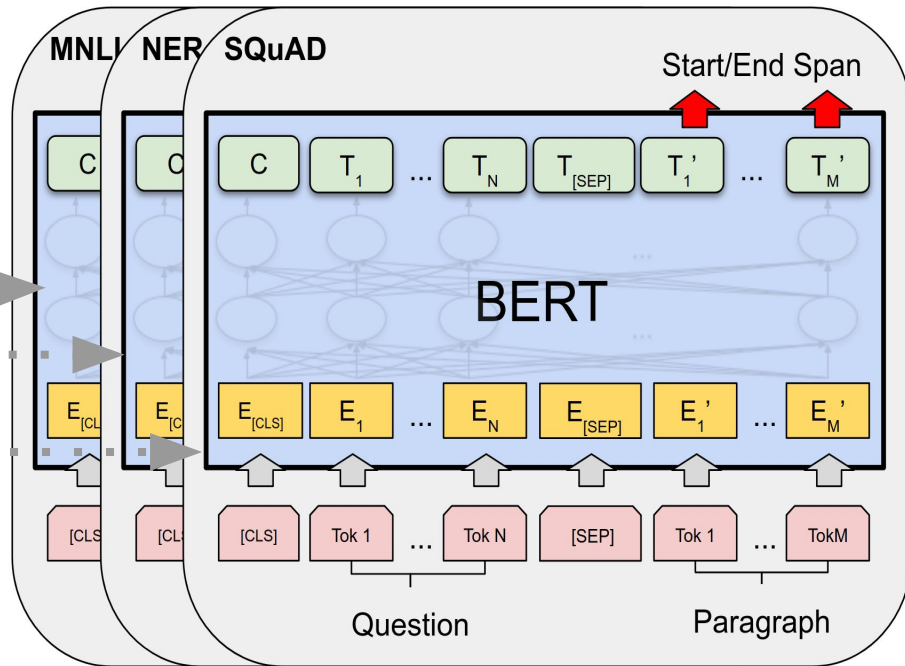


# BERT: Transformers + Pre-training + Fine-Tuning

## Pre-Training



## Fine-Tuning



# Transformers + Pre-training -- new dawn of NLP



System	MNLI-(m/mm) 392k	QQP 363k	QNLI 108k	SST-2 67k	CoLA 8.5k	STS-B 5.7k	MRPC 3.5k	RTE 2.5k	Average -
Pre-OpenAI SOTA	80.6/80.1	66.1	82.3	93.2	35.0	81.0	86.0	61.7	74.0
BiLSTM+ELMo+Attn	76.4/76.1	64.8	79.8	90.4	36.0	73.3	84.9	56.8	71.0
OpenAI GPT	82.1/81.4	70.3	87.4	91.3	45.4	80.0	82.3	56.0	75.1
BERT <sub>BASE</sub>	84.6/83.4	71.2	90.5	93.5	52.1	85.8	88.9	66.4	79.6
BERT <sub>LARGE</sub>	<b>86.7/85.9</b>	<b>72.1</b>	<b>92.7</b>	<b>94.9</b>	<b>60.5</b>	<b>86.5</b>	<b>89.3</b>	<b>70.1</b>	<b>82.1</b>

Devlin et al. 2019

Yang et al. 2019

Lan et al. 2019

# Transformers + Pre-training -- new dawn of NLP

System	MNLI-(m/mm) 392k	QQP 363k	QNLI 108k	Dataset	XLNet-Large (as in paper)	XLNet-Large -wikibooks	BERT-Large -wikibooks	Average
							best of 3 variants	
Pre-OpenAI SOTA	80.6/80.1	66.1	82.3	SQuAD1.1 EM	89.0	88.2	86.7 (II)	-
BiLSTM+ELMo+Attn	76.4/76.1	64.8	79.8	SQuAD1.1 F1	94.5	94.0	92.8 (II)	4.0
OpenAI GPT	82.1/81.4	70.3	87.4	SQuAD2.0 EM	86.1	85.1	82.8 (II)	1.0
BERT <sub>BASE</sub>	84.6/83.4	71.2	90.5	SQuAD2.0 F1	88.8	87.8	85.5 (II)	5.1
BERT <sub>LARGE</sub>	<b>86.7/85.9</b>	<b>72.1</b>	<b>92.7</b>	RACE	81.8	77.4	75.1 (II)	9.6
				MNLI	89.8	88.4	87.3 (II)	2.1
				QNLI	93.9	93.9	93.0 (II)	
				QQP	91.8	91.8	91.4 (II)	
				RTE	83.8	81.2	74.0 (III)	
				SST-2	95.6	94.4	94.0 (II)	
				MRPC	89.2	90.0	88.7 (III)	
				CoLA	63.6	65.2	63.7 (II)	
				STS-B	91.8	91.1	90.2 (III)	

Devlin et al. 2019

Yang et al. 2019

Lan et al. 2019

# Transformers + Pre-training -- new dawn of NLP

				Dataset		XLNet-Large (as in paper)		XLNet-Large -wikibooks		BERT-Large -wikibooks		Average	
System	MNLI-(m/mm)	QQP	QNLI	best of 3 variants									
	392k	363k	108k	SQuAD1.1 EM	89.0	88.2	86.7 (II)						-
Pre-OpenAI SOTA	80.6/80.1	66.1	82.3	SQuAD1.1 F1	94.5	94.0	92.8 (II)						4.0
BiLSTM+ELMo+Attn	76.4/76.1	64.8	79.8	SQuAD2.0 EM	86.1	85.1	82.8 (II)						1.0
OpenAI GPT	82.1/81.4	70.3	87.4	SQuAD2.0 F1	88.8	87.8	85.5 (II)						5.1
BERT <sub>BASE</sub>	Models	MNLI	QNLI	QQP	RTE	SST	MRPC	CoLA	STS	WNLI	Avg	9.6	
BERT <sub>LARGE</sub>	Single-task single models on dev											2.1	
	BERT-large	86.6	92.3	91.3	70.4	93.2	88.0	60.6	90.0	-	-		
	XLNet-large	89.8	93.9	91.8	83.8	95.6	89.2	63.6	91.8	-	-		
	RoBERTa-large	90.2	94.7	92.2	86.6	96.4	90.9	68.0	92.4	-	-		
	ALBERT (1M)	90.4	95.2	92.0	88.1	96.8	90.2	68.7	92.7	-	-		
	ALBERT (1.5M)	90.8	95.3	92.2	89.2	96.9	90.9	71.4	93.0	-	-		
	Ensembles on test (from leaderboard as of Sept. 16, 2019)												
	ALICE	88.2	95.7	90.7	83.5	95.2	92.6	69.2	91.1	80.8	87.0		
	MT-DNN	87.9	96.0	89.9	86.3	96.5	92.7	68.4	91.1	89.0	87.6		
	XLNet	90.2	98.6	90.3	86.3	96.8	93.0	67.8	91.6	90.4	88.4		
	RoBERTa	90.8	98.9	90.2	88.2	96.7	92.3	67.8	92.2	89.0	88.5		
	Adv-RoBERTa	91.1	98.8	90.3	88.7	96.8	93.1	68.0	92.4	89.0	88.8		
	ALBERT	91.3	99.2	90.5	89.2	97.1	93.4	69.1	92.5	91.8	89.4		

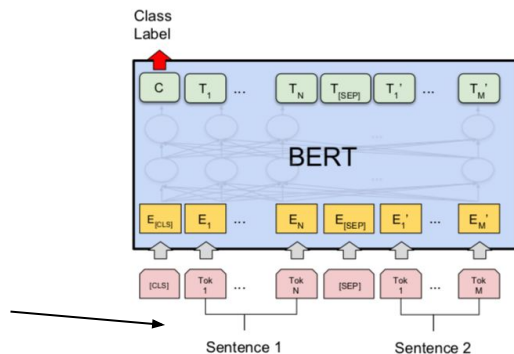
Devlin et al. 2019

Yang et al. 2019

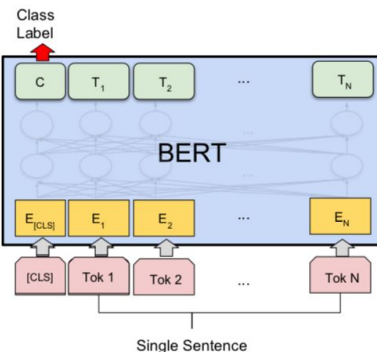
Lan et al. 2019

# BERT: Fine-tuning Paradigms

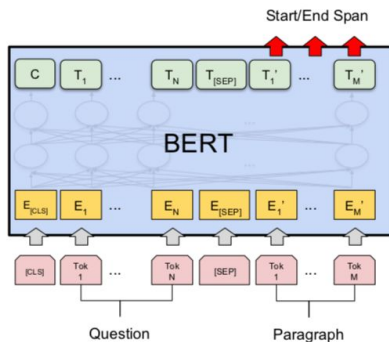
This can be a  
relevance  
scoring model



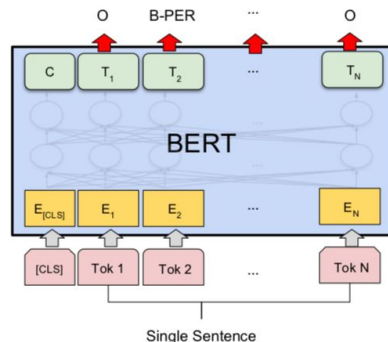
(a) Sentence Pair Classification Tasks:  
MNLI, QQP, QNLI, STS-B, MRPC,  
RTE, SWAG



(b) Single Sentence Classification Tasks:  
SST-2, CoLA



(c) Question Answering Tasks:  
SQuAD v1.1



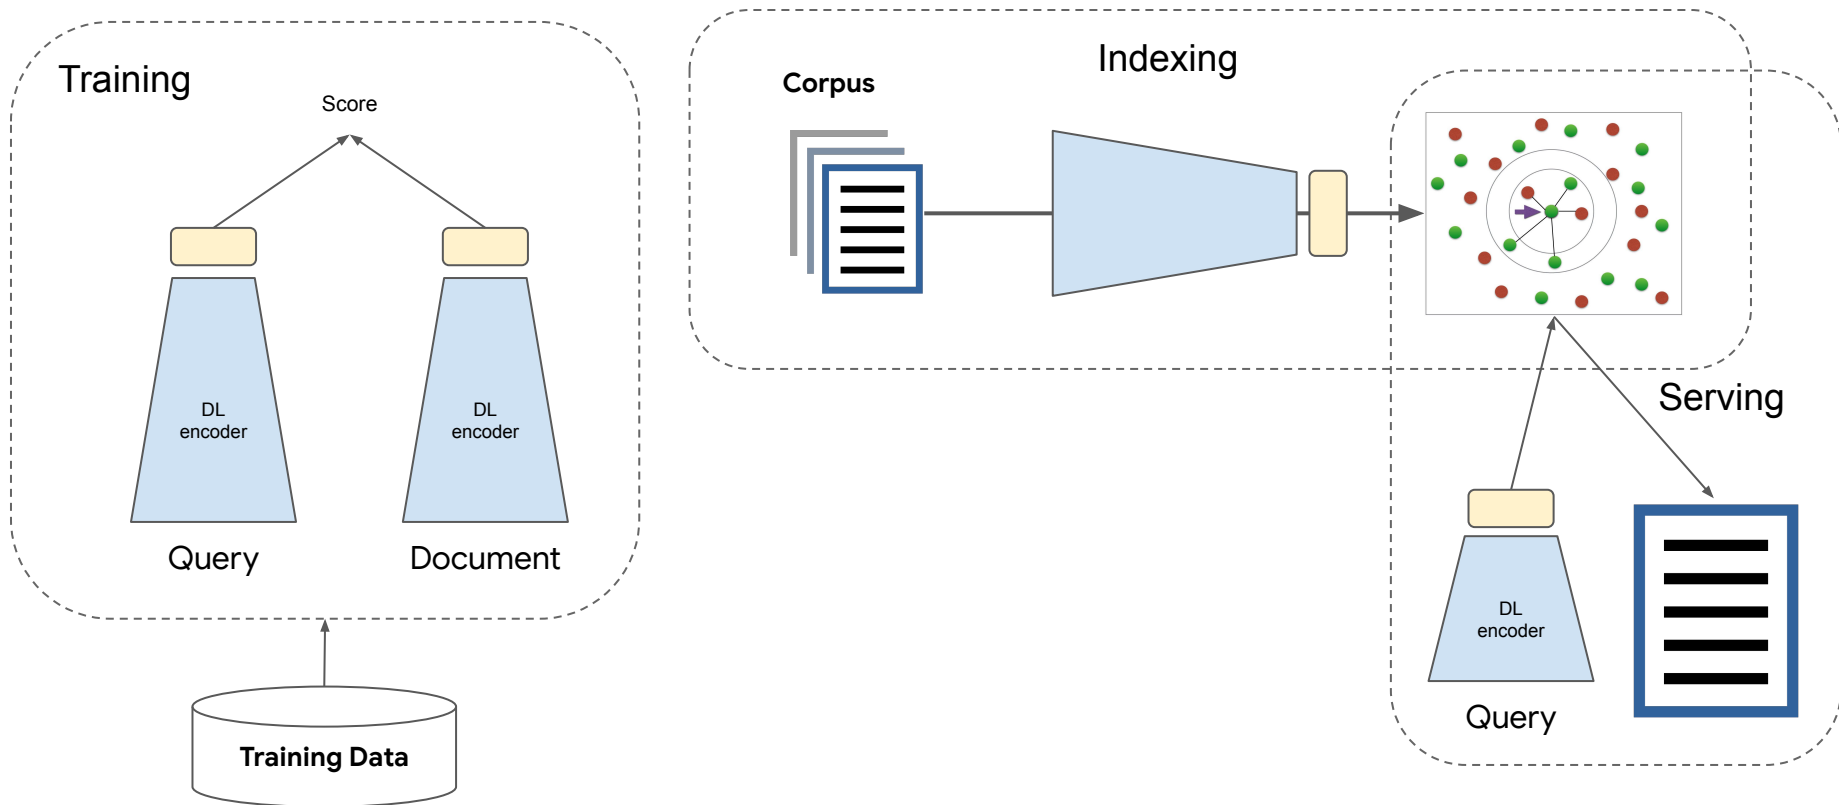
(d) Single Sentence Tagging Tasks:  
CoNLL-2003 NER

# BERT 4 Document Relevance Scoring

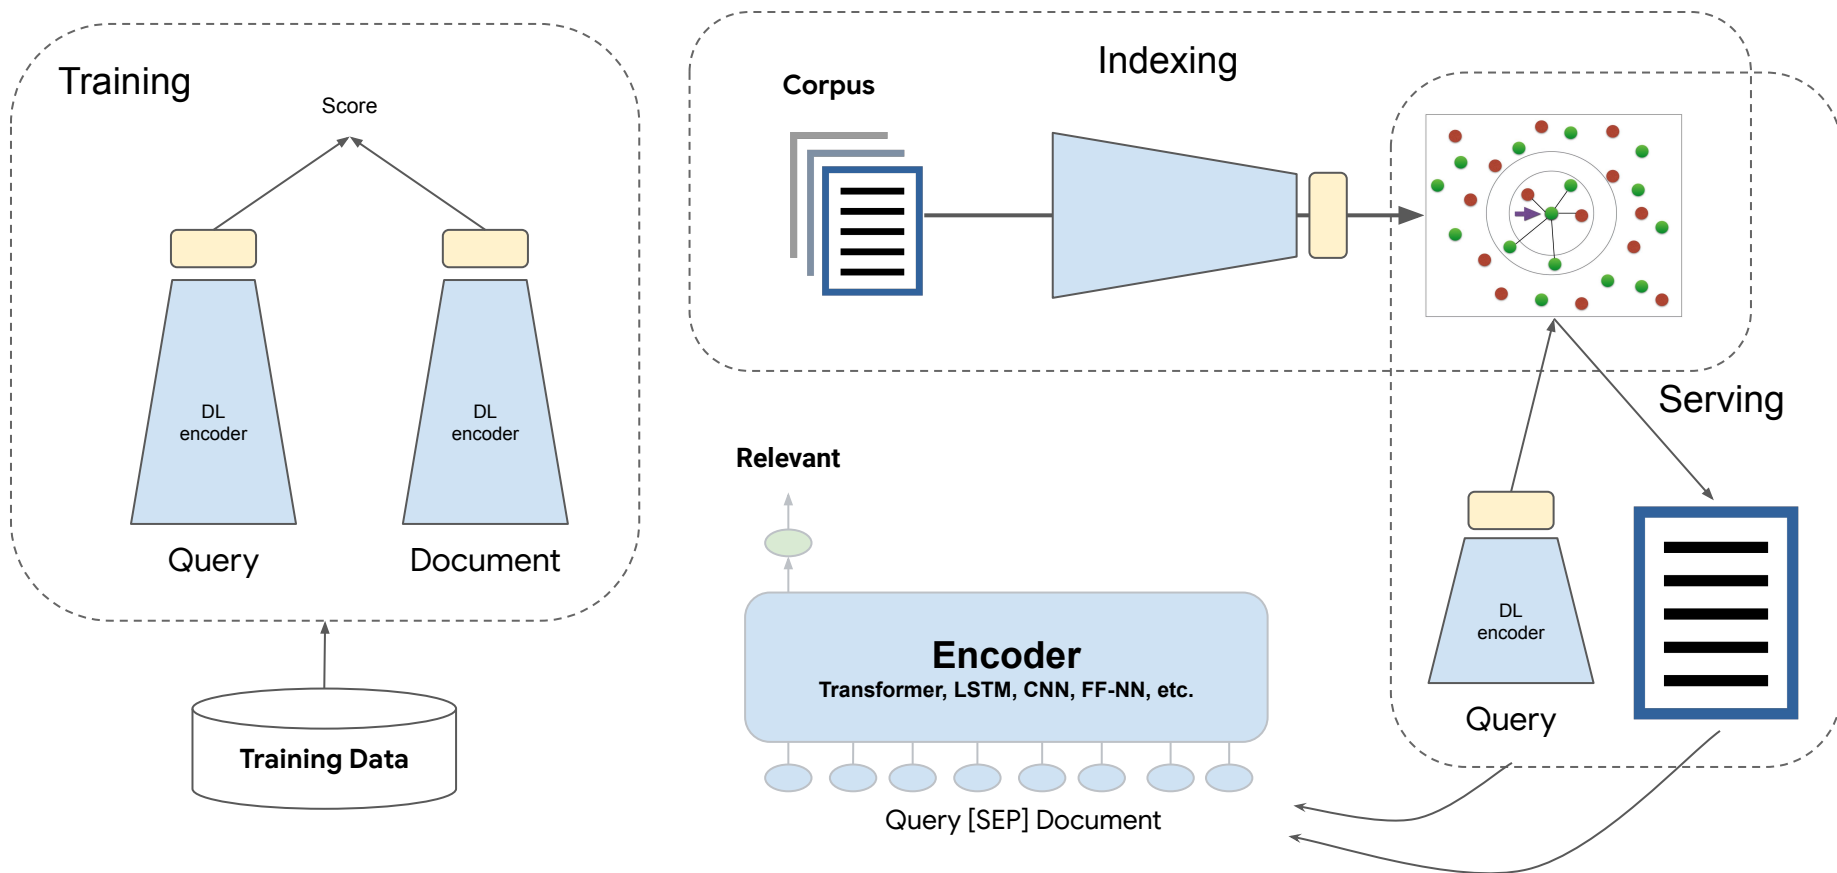
CEDR: Macaveney et al. 2019

Ranker	Input Representation	Robust04		WebTrack 2012-14	
		P@20	nDCG@20	nDCG@20	ERR@20
BM25	n/a	0.3123	0.4140	0.1970	0.1472
SDM [13]	n/a	0.3749	0.4353	-	-
TREC-Best	n/a	<b>0.4386</b>	<b>0.5030</b>	0.2855	<b>0.2530</b>
ConvKNRM	GloVe	0.3349	0.3806	[B] 0.2547	[B] 0.1833
Vanilla BERT	BERT (fine-tuned)	[BC] 0.4042	[BC] 0.4541	<b>[BC] 0.2895</b>	[BC] 0.2218
PACRR	GloVe	0.3535	[C] 0.4043	0.2101	0.1608
PACRR	ELMo	[C] 0.3554	[C] 0.4101	[BG] 0.2324	[BG] 0.1885
PACRR	BERT	[C] 0.3650	[C] 0.4200	0.2225	0.1817
PACRR	BERT (fine-tuned)	[BCVG] 0.4492	[BCVG] 0.5135	[BCG] 0.3080	[BCG] 0.2334
CEDR-PACRR	BERT (fine-tuned)	<b>[BCVG] 0.4559</b>	<b>[BCVG] 0.5150</b>	<b>[BCVGN] 0.3373</b>	<b>[BCVGN] 0.2656</b>
KNRM	GloVe	0.3408	0.3871	[B] 0.2448	0.1755
KNRM	ELMo	[C] 0.3517	[CG] 0.4089	0.2227	0.1689
KNRM	BERT	[BCG] 0.3817	[CG] 0.4318	[B] 0.2525	[B] 0.1944
KNRM	BERT (fine-tuned)	[BCG] 0.4221	[BCVG] 0.4858	[BCVG] 0.3287	[BCVG] 0.2557
CEDR-KNRM	BERT (fine-tuned)	<b>[BCVGN] 0.4667</b>	<b>[BCVGN] 0.5381</b>	<b>[BCVG] 0.3469</b>	<b>[BCVG] 0.2772</b>
DRMM	GloVe	0.2892	0.3040	0.2215	0.1603
DRMM	ELMo	0.2867	0.3137	[B] 0.2271	0.1762
DRMM	BERT	0.2878	0.3194	[BG] 0.2459	[BG] 0.1977
DRMM	BERT (fine-tuned)	[CG] 0.3641	[CG] 0.4135	[BG] 0.2598	[B] 0.1856
CEDR-DRMM	BERT (fine-tuned)	<b>[BCVGN] 0.4587</b>	<b>[BCVGN] 0.5259</b>	<b>[BCVGN] 0.3497</b>	<b>[BCVGN] 0.2621</b>

# Dual Encoder Retrieval (AKA two-tower; AKA relevance-based)

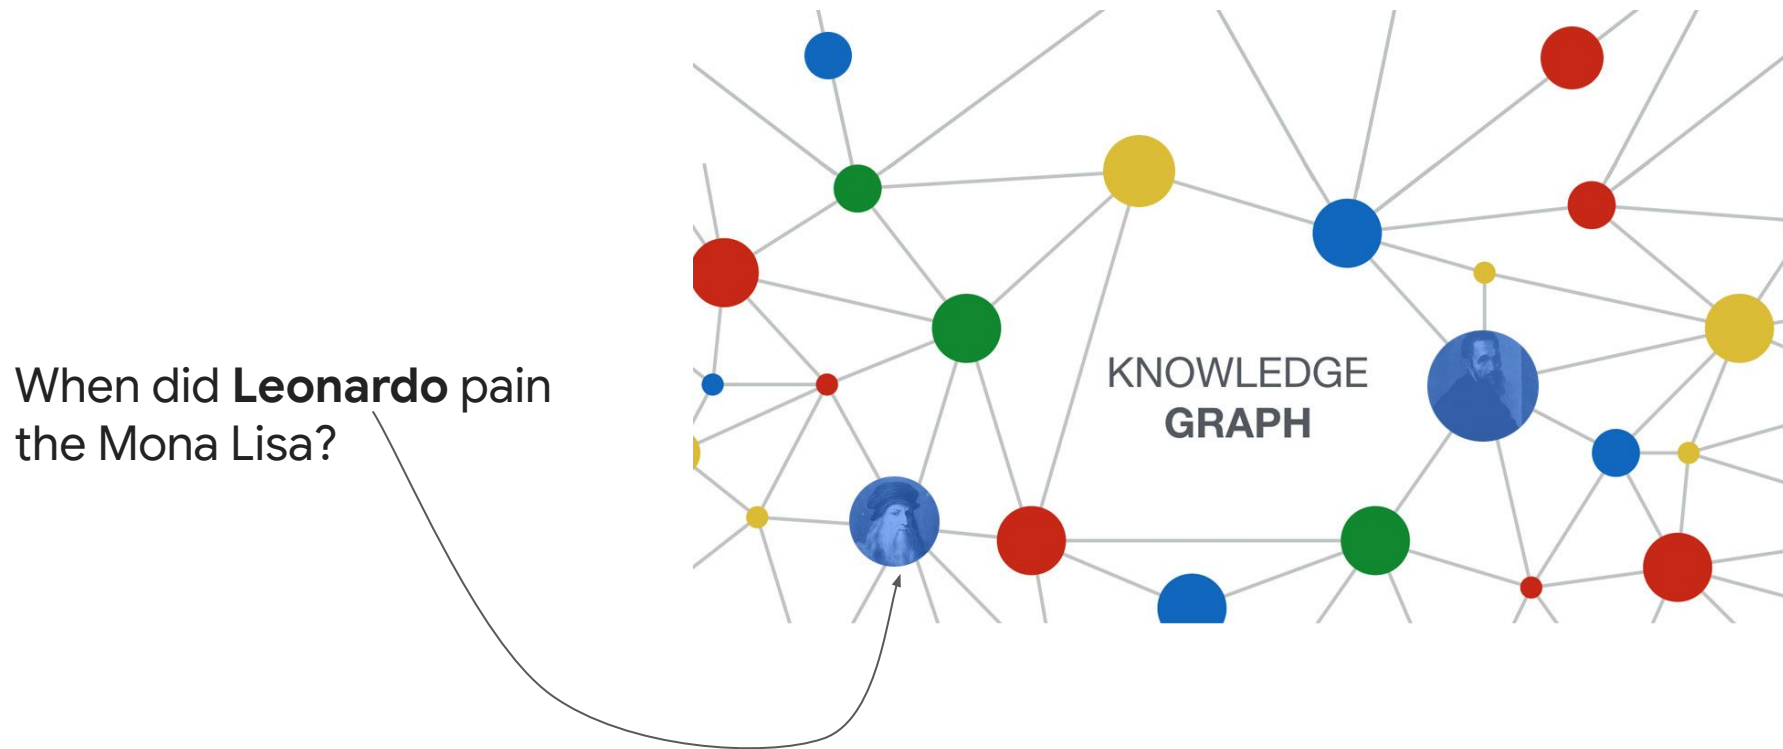


# Dual Encoder Retrieval (AKA two-tower; AKA relevance-based)






# Classification with massive output spaces



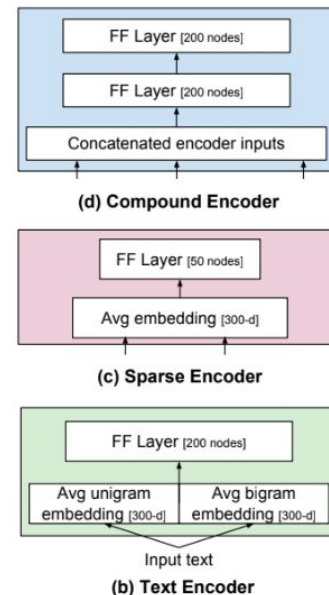
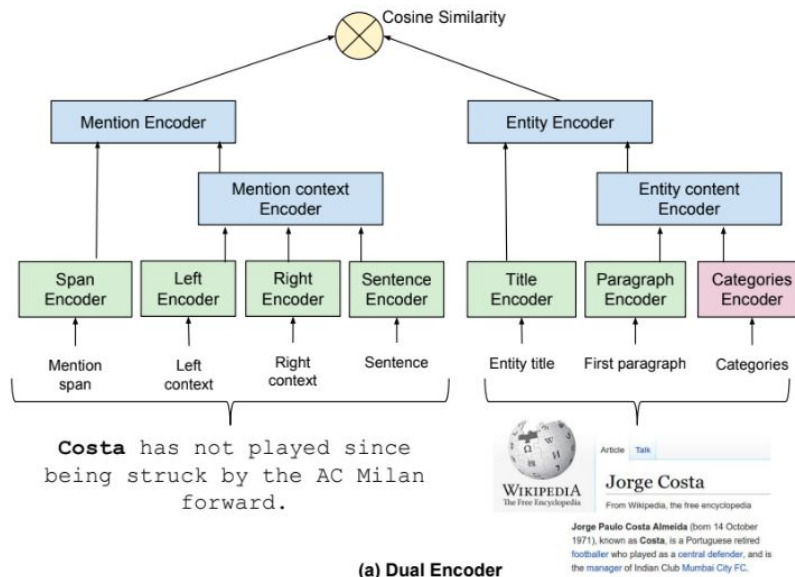
# Embedding **Objects** vs. Descriptions



# Embedding Objects vs. Descriptions

Dataset	AT-Prior	R@1	
		DE-RN	DE-MN
CoNLL	65.71	40.87	<b>77.93</b>
ACE2004	80.93	55.64	<b>87.55</b>
AQUAINT	82.64	54.55	<b>86.78</b>
MSNBC	65.62	42.11	<b>75.30</b>
WikiSample	79.01	59.17	<b>84.06</b>
TACKBP 2009	69.40	51.34	<b>78.60</b>
TACKBP 2010	72.45	48.63	<b>87.35</b>
TACKBP 2011	55.52	35.85	<b>73.04</b>
TACKBP 2012	26.45	21.85	<b>49.91</b>
Wikinews	86.66	66.25	<b>91.56</b>
Average	68.44	47.63	<b>79.21</b>

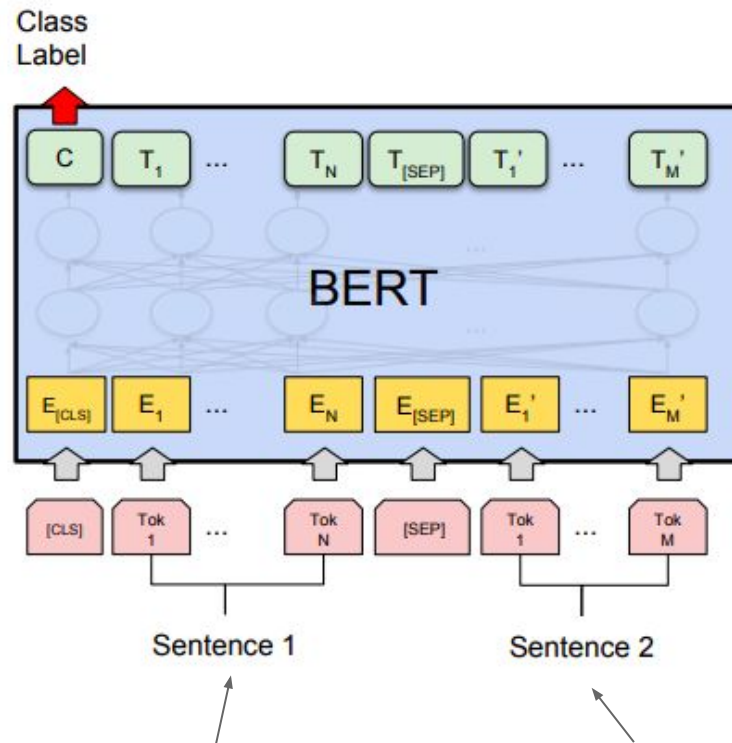
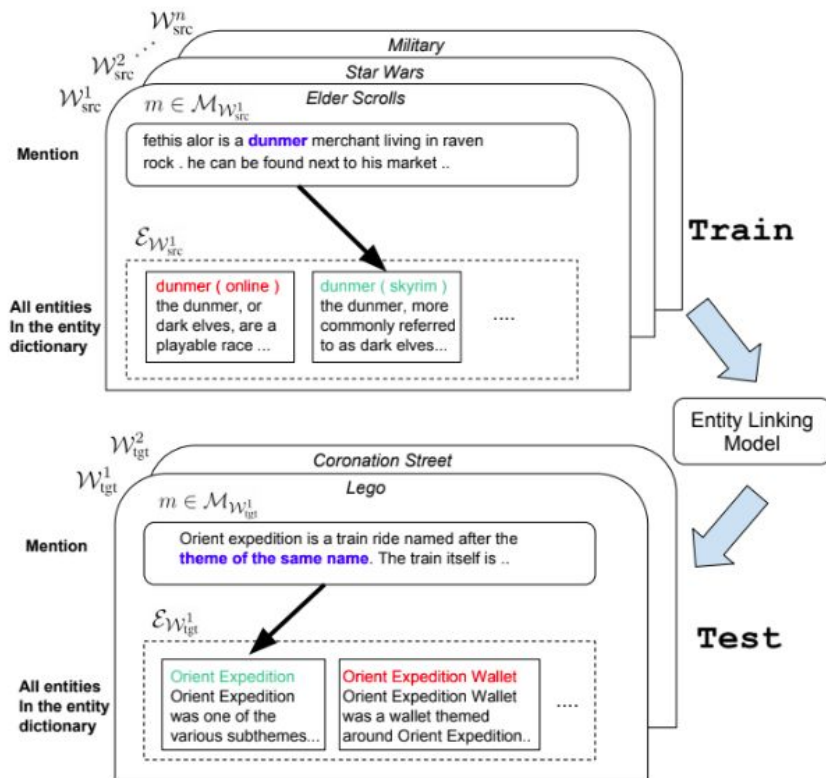
Gillick et al. 2019



$f(\text{text}, \text{text})$

# Zero-Shot Entity Linking

Lee et al. 2019




Fethis alor is a **dunmer\*** merchant living in raven rock.

\* special mention vector added to indicate this is the focus mention

**dunmer ( skyrim )** the dunmer, more commonly referred to as dark elves

# Zero-Shot Entity Linking


Lee et al. 2019



Pretraining	$\mathcal{W}_{tgt}^1$	$\mathcal{W}_{tgt}^2$	$\mathcal{W}_{tgt}^3$	$\mathcal{W}_{tgt}^4$	Avg
$U_{src+tgt}$ (Glorot et al., 2011) <sup>†</sup>	73.19	71.61	62.16	64.69	67.91
$U_{src+tgt} \rightarrow U_{tgt}$ (DAP)	79.20	75.55	66.85	66.72	<b>72.08</b>
$U_{WB}$ (Devlin et al., 2019)	83.40	79.00	73.03	68.82	76.06
$U_{WB} \rightarrow U_{tgt}$ (DAP)	81.68	81.34	73.17	71.97	<b>77.04</b>
$U_{WB} \rightarrow U_{src+tgt}$	82.92	79.00	72.62	69.55	76.02
$U_{WB} \rightarrow U_{src+tgt} \rightarrow U_{tgt}$ (DAP)	82.82	81.59	75.34	72.52	<b>78.07</b>

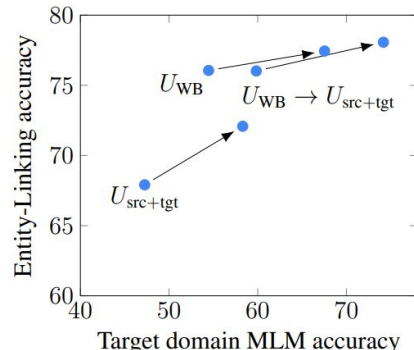
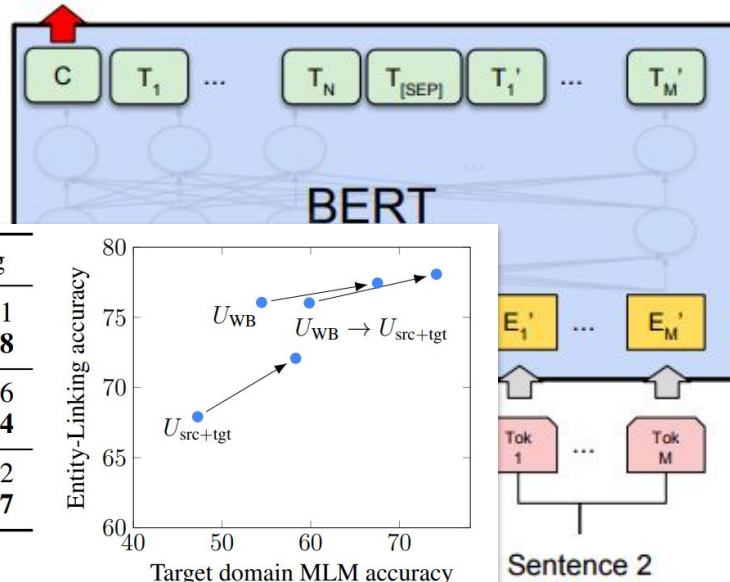
Mention

All entities  
in the entity  
dictionary



Test

Class  
Label



Fethis alor is a **dunmer\*** merchant living in raven rock.


\* special mention vector added to indicate this is the focus mention

**dunmer ( skyrim )** the dunmer, more commonly referred to as dark elves

# Knowledge Retrieval → Comprehension



why is the sky blue



[All](#) [Videos](#) [Images](#) [Shopping](#) [Books](#) [More](#) [Settings](#) [Tools](#)

About 3,490,000,000 results (0.53 seconds)

The Short Answer:

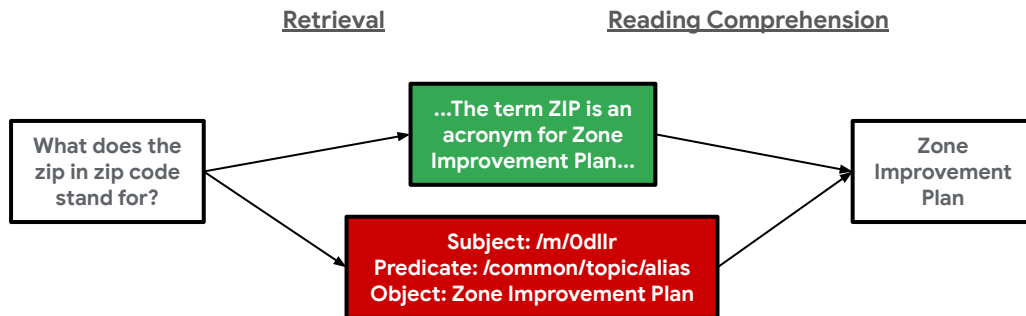
Sunlight reaches Earth's atmosphere and is scattered in all directions by all the gases and particles in the air. **Blue** light is scattered more than the other colors because it travels as shorter, smaller waves. This is why we see a **blue sky** most of the time.

[Why Is the Sky Blue? | NASA Space Place – NASA Science for ...](https://spaceplace.nasa.gov/blue-sky)  
<https://spaceplace.nasa.gov/blue-sky>

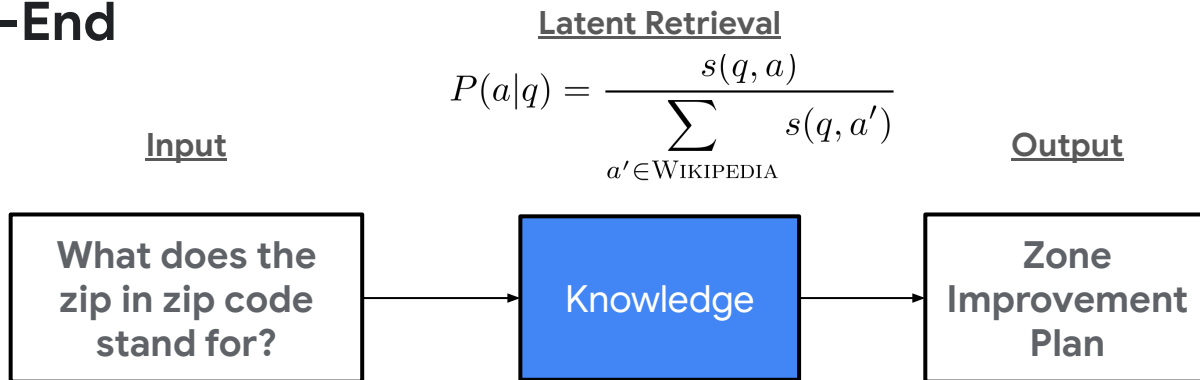
[About](#) [Featured Snippets](#) [Feedback](#)

Reading/QA

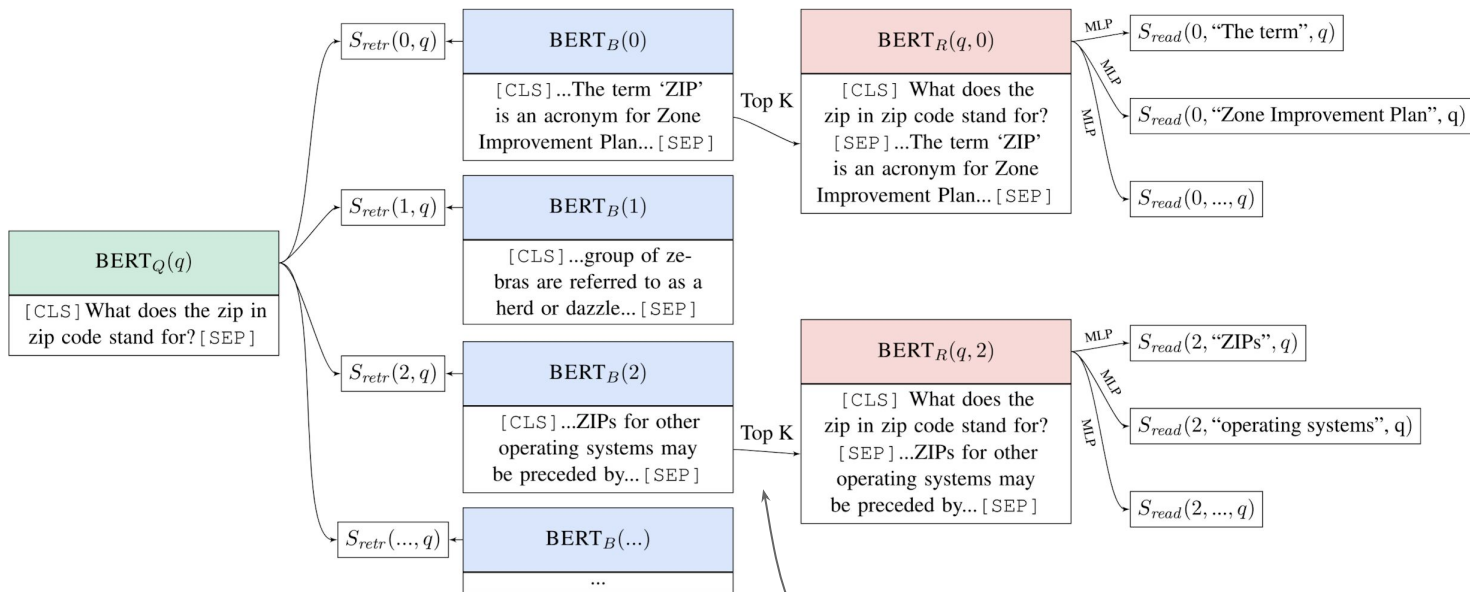
# Pipelined Approach



# End-to-End



# ORQA Overview




$$P(a|q) = \frac{s(q, a)}{\sum_{a' \in \text{WIKIPEDIA}} s(q, a')}$$

Approx w/ top-K



# Results

End-to-end learning is crucial for information-seeking queries!



## Sequential Question Answering

Building	City	Floors
First Canadian Place	Toronto	72
Commerce Court West	Toronto	57
Tour de la Bourse	Montreal	47
Place Ville-Marie	Montreal	44

**What are the buildings in Toronto?**  
*First Canadian Place, Commerce Court West*

## Sequential Question Answering

Building	City	Floors
First Canadian Place	Toronto	72
Commerce Court West	Toronto	57
Tour de la Bourse	Montreal	47
Place Ville-Marie	Montreal	44

**What are the buildings in Toronto?**

*First Canadian Place, Commerce Court West*

**Of those, which buildings have more than 60 floors?**

*First Canadian Place*

## Sequential Question Answering

Building	City	Floors
First Canadian Place	Toronto	72
Commerce Court West	Toronto	57
Tour de la Bourse	Montreal	47
Place Ville-Marie	Montreal	44

**What are the buildings in Toronto?**

*First Canadian Place, Commerce Court West*

**Of those, which buildings have more than 60 floors?**

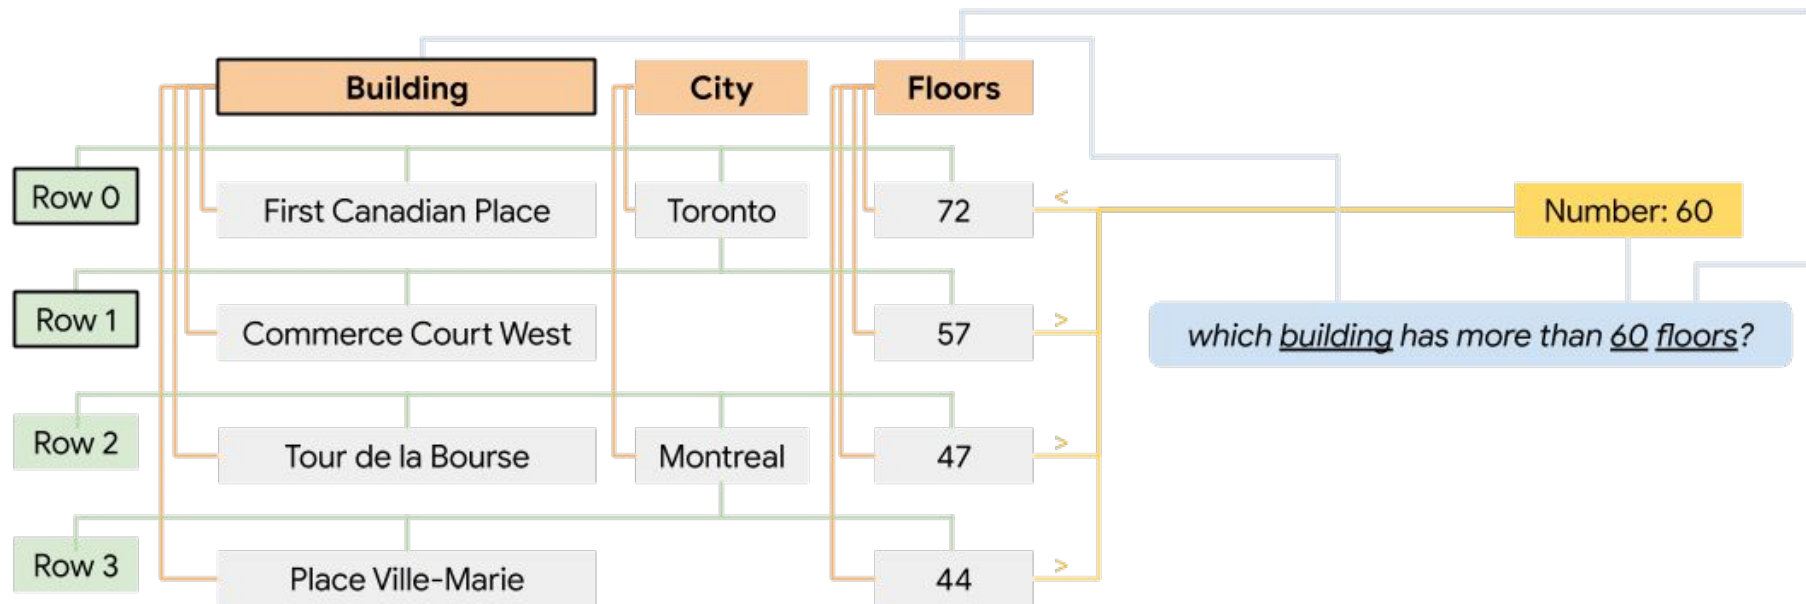
*First Canadian Place*

**How many floors does it have?**

*72*

# Graph Transformer

Müller et al. 2019



## Results on SQA

Model	ALL	Q1	Q2	Q3
Iyyer et al. (2017) <sup>†</sup>	44.7	70.4	41.1	23.6
Sun et al. (2018) <sup>†</sup> *	45.6	70.3	42.6	24.8
Müller et al. (2019) <sup>†</sup> *	55.1	67.2	52.7	46.8

# Retrieval + NLP

