End-to-End Training of Multi-Document Reader and Retriever for Open-Domain Question Answering

Devendra Singh Sachan[1,2], Siva Reddy[1,2], Will Hamilton[1,2], Chris Dyer[3], Dani Yogatama[3]

[1]McGill University, [2]Mila, [3]DeepMind

sachande@mila.quebec, dyogatama@deepmind.com

Introduction

Open-Domain Question Answering

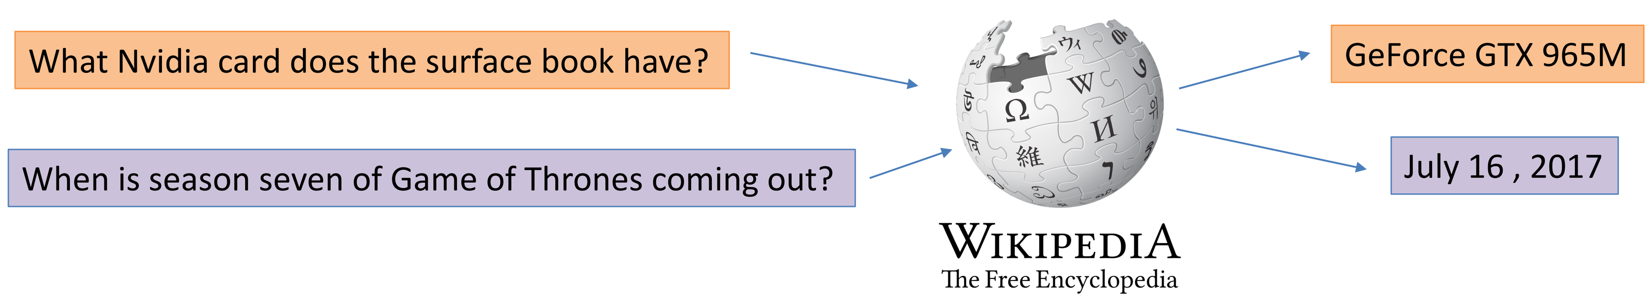

- **Input**: Question (q) and evidence documents (D) such as Wikipedia
- **Output**: Answer (a)

Retriever: Dual Encoder

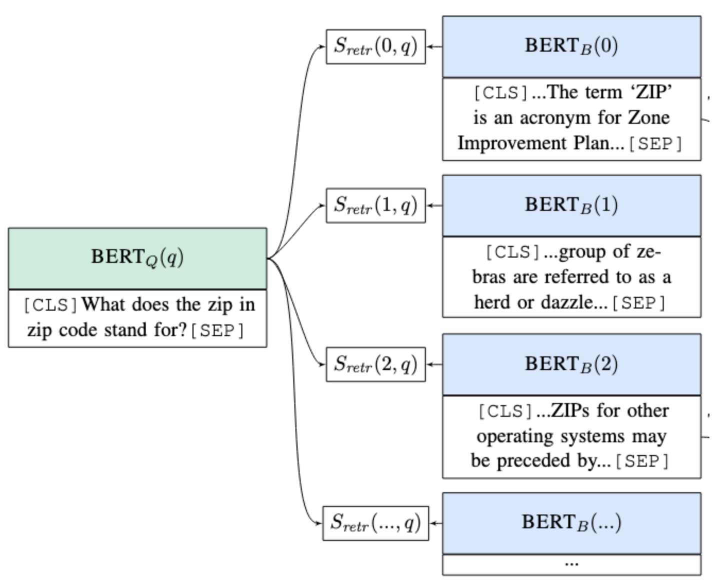

- Evidence Documents

$$\mathcal{D} = \{\boldsymbol{d}_1, \ldots, \boldsymbol{d}_M\}$$

$$\mathrm{score}(\boldsymbol{q}, \boldsymbol{d}_i; \Phi) = f_q(\boldsymbol{q}; \Phi_q)^\top f_d(\boldsymbol{d}_i; \Phi_d)$$

- Select Top-K Documents with highest scores

$$\mathcal{Z} = \{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_K\}$$

Multi-Document Reader: Fusion-in-Decoder (FiD)

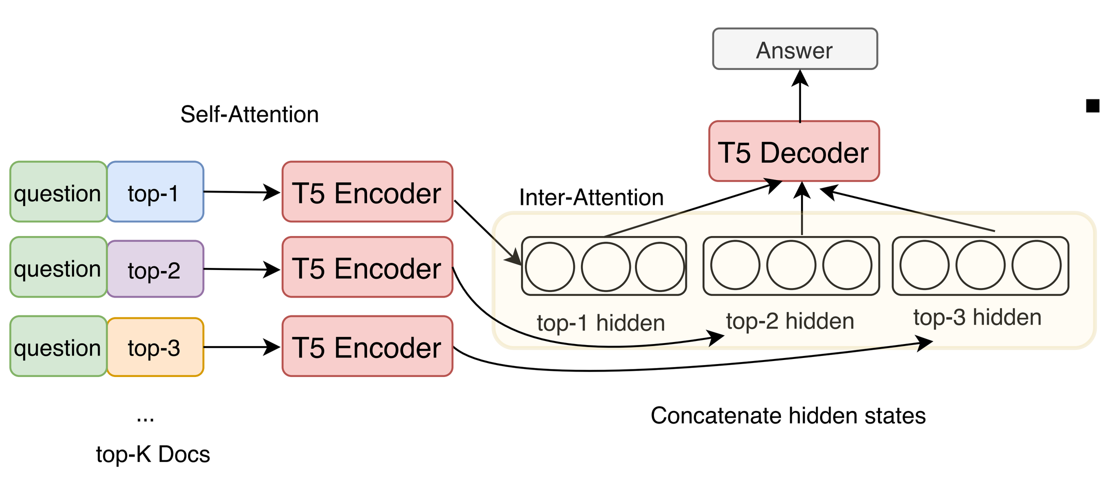

- Independent self-attention
- Concatenate encoder representations for decoder's inter-attention
- Autoregressive training

Research Question

How to jointly train the FiD reader and dual-encoder retriever?

Methods

EMDR^2: Training Objective

$$\mathcal{L} = \underbrace{\log p(\boldsymbol{a} \mid \boldsymbol{q}, \mathcal{Z}_{\text{top-}K}; \Theta)}_{\text{reader}} + \underbrace{\log \sum_{k=1}^{K} \mathbb{SG}\left(p(\boldsymbol{a} \mid \boldsymbol{q}, \boldsymbol{z}_k; \Theta)\right) p(\boldsymbol{z}_k \mid \boldsymbol{q}, \mathcal{Z}_{\text{top-}K}; \Phi)}_{\text{retriever}},$$

Training Pipeline

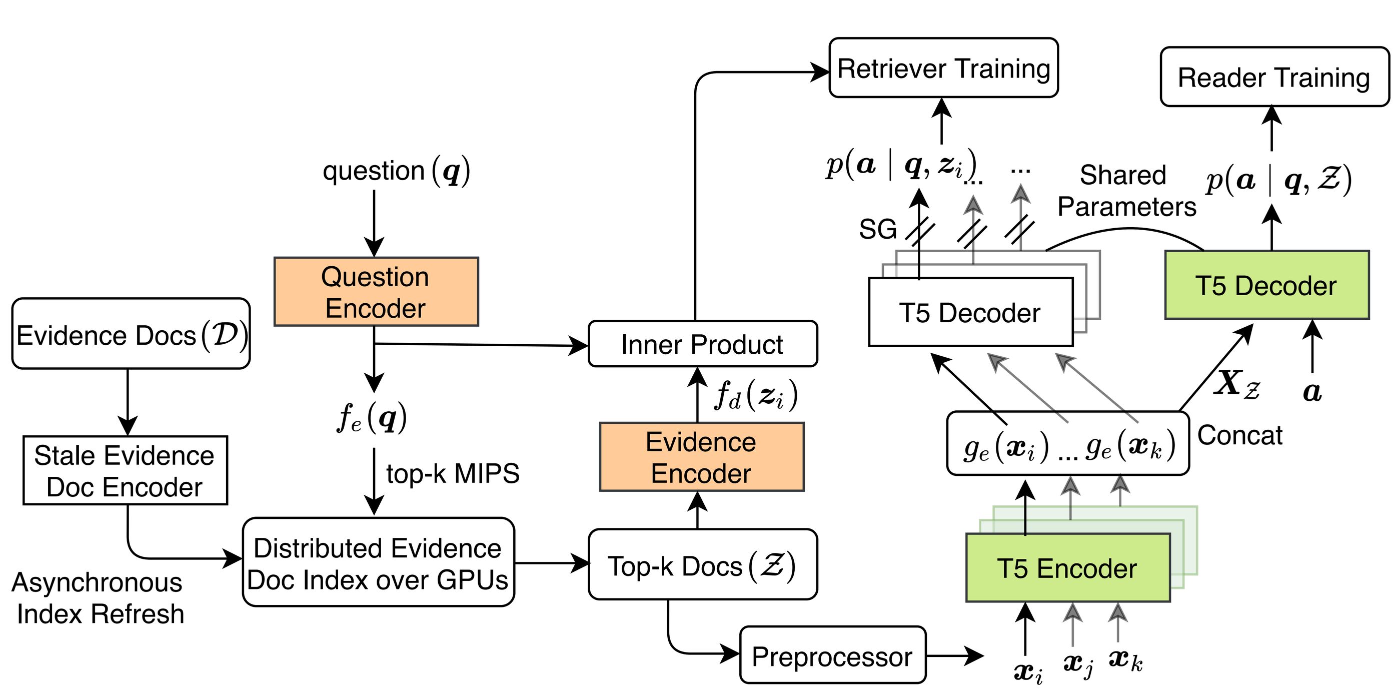

EMDR^2: Expectation-Maximization View

Algorithm 1: End-to-end training of multi-document reader and retriever.

Input: Model parameters Θ and Φ, evidence documents $\mathcal{D}$.
while *not converged* **do**
- Compute $\mathcal{Z}_{\text{top-}K}$ using the current retriever parameters Φ. `// E-step`
- Compute $p(\boldsymbol{a} \mid \boldsymbol{q}, \boldsymbol{z}_k)$ for each $\boldsymbol{z}_k$ using the current reader parameters Θ. `// E-step`
- Update model parameters Θ and Φ to maximize the log-likelihood in Eq. 6. `// M-step`

end

Comparison with Recent Methods

Methods	Multi-Doc Reader	Retriever Adaptation	End-to-End Training	Unsupervised Retriever
REALM		✓	✓	✓
DPR				
RAG		✓	✓	
FiD	✓			
FiD-KD	✓	✓		
EMDR²	✓	✓	✓	✓

Results and Analysis

- Bootstrap the model with unsupervised **masked salient spans (MSS)** training.

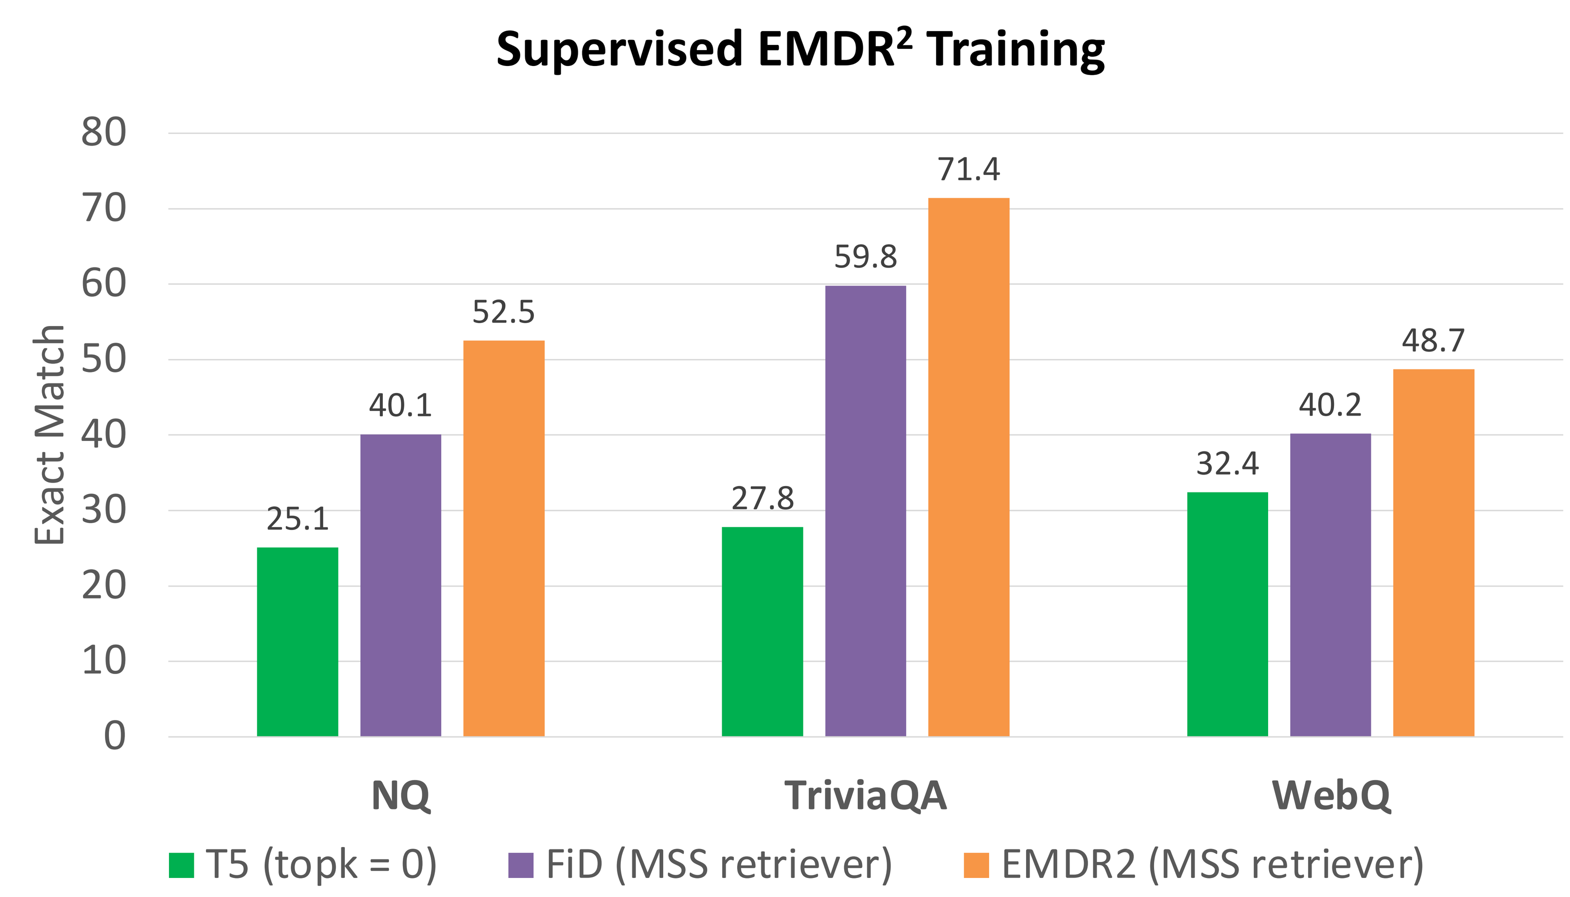

- 8-12 EM points gain over no end-to-end training.

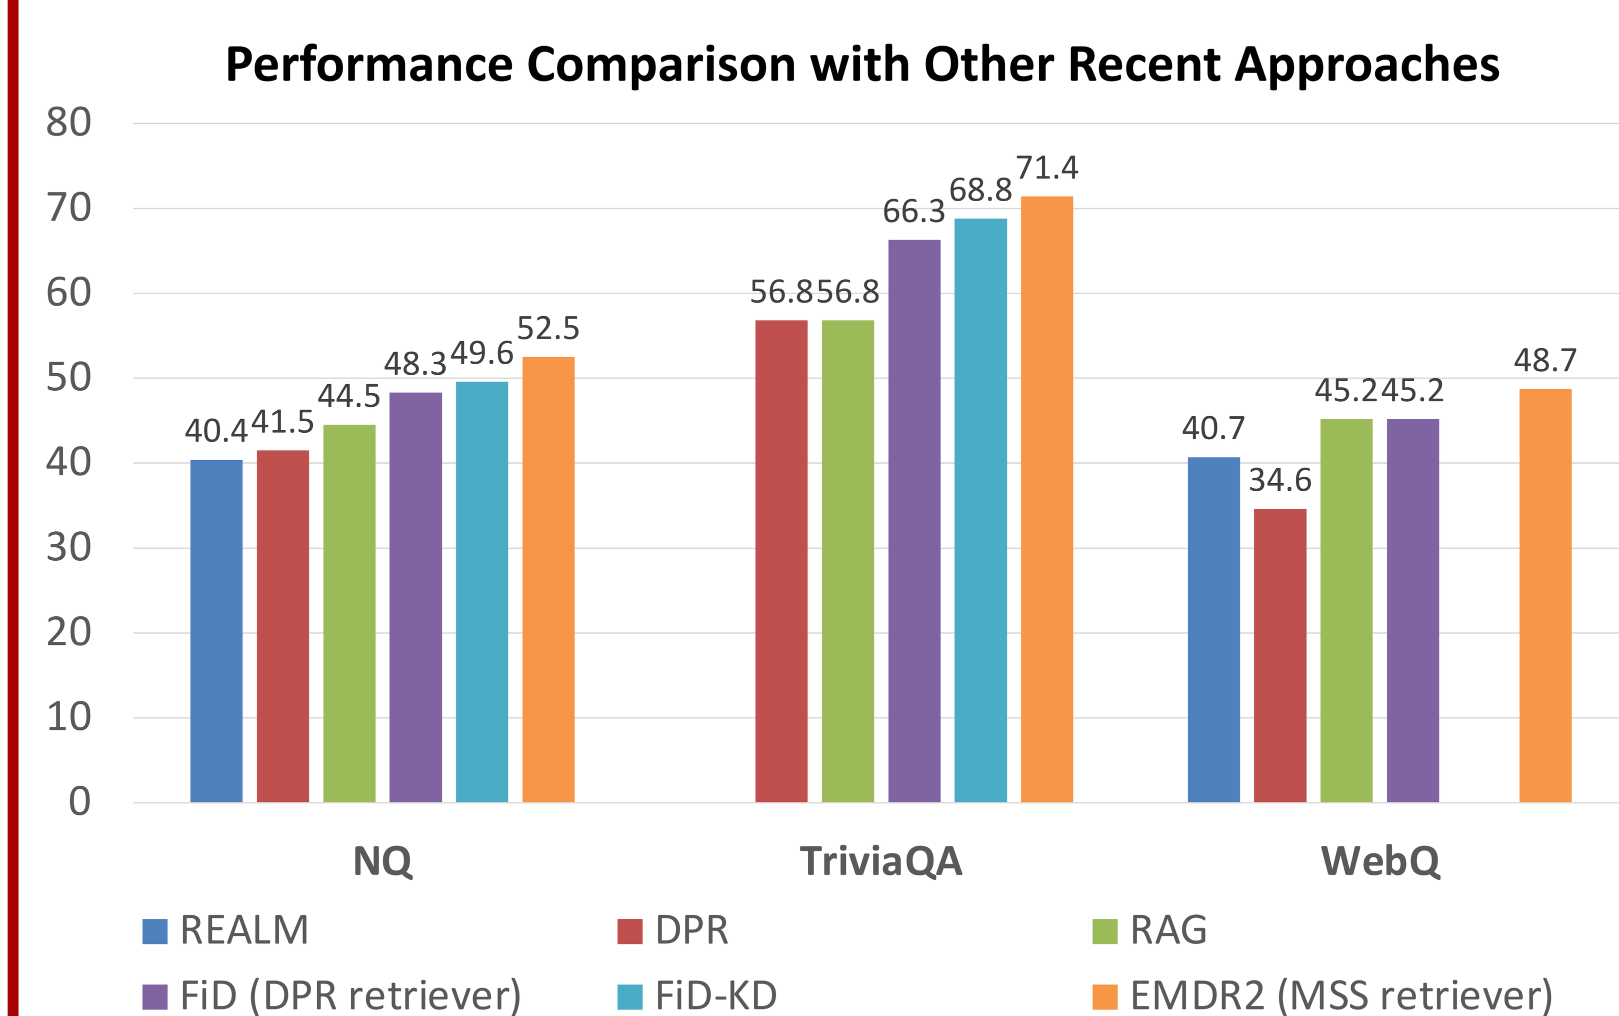

- 2-3 EM points gain over previous SOTA results.

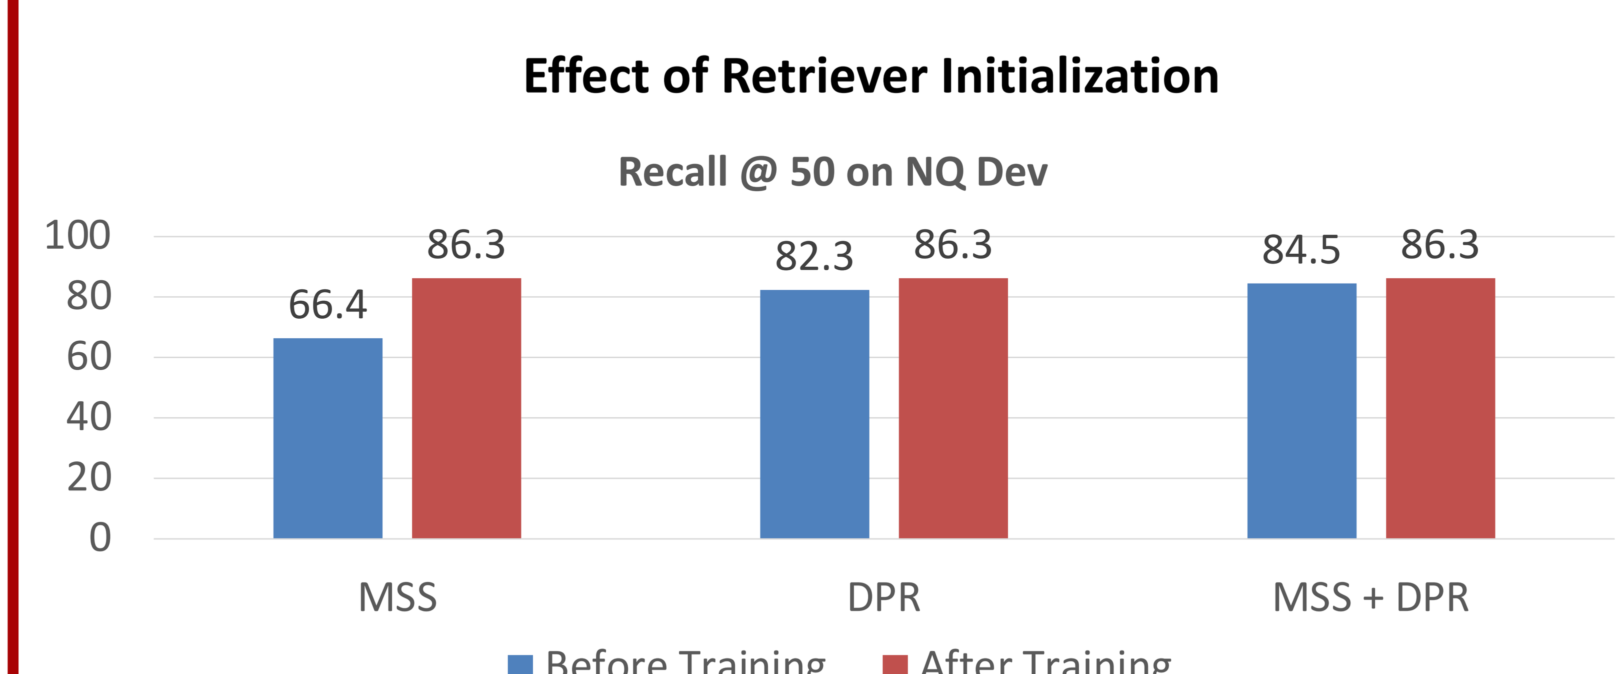

- Different retriever initializations converge to the same final recall.
- Additional retriever training by DPR shows no further performance gains.

Conclusions

- We obtain *state-of-the-art results* with a new end-to-end training algorithm.
- Requires *single training cycle* of retriever and reader training.
- *Supervised retriever* initialization may not be necessary for SOTA performance.

NeurIPS 2021

Code: https://github.com/DevSinghSachan/emdr2