Meta-Learning Neural Bloom Filters

Jack Rae Sergey Bartunov Tim Lillicrap



DeepMind

Architecture

Interested in neural networks with compressive, distributed memories.

Problem

Trend in the use of neural networks to replace classical data-structures.

Bloom Filter

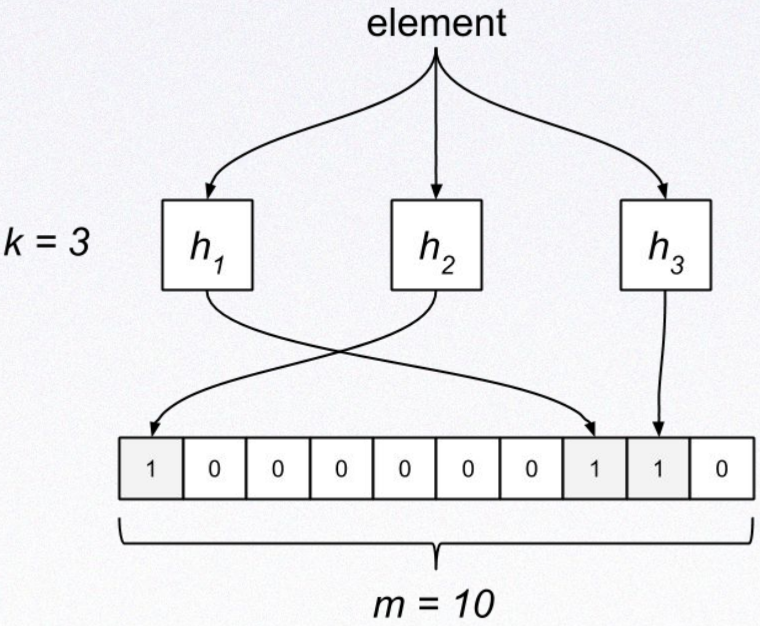

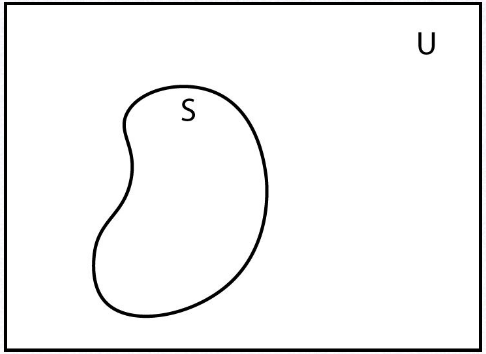

Bloom Filter

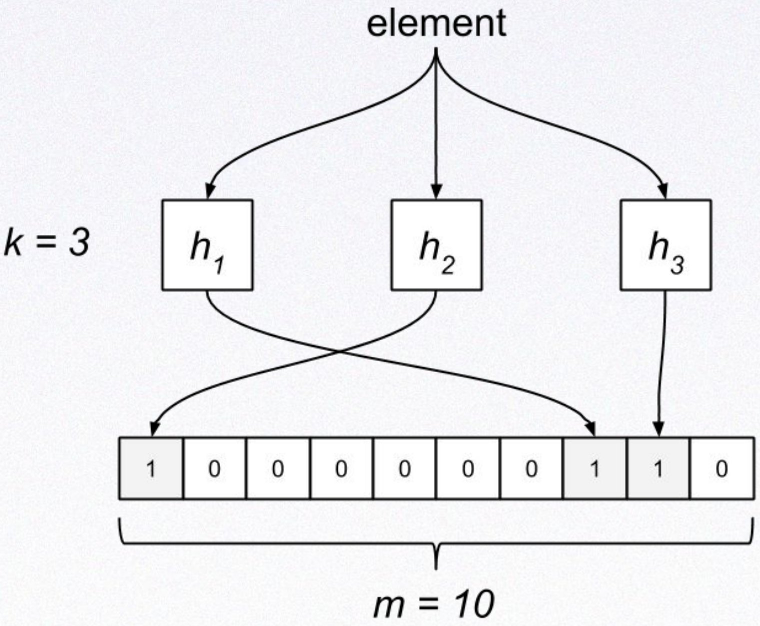

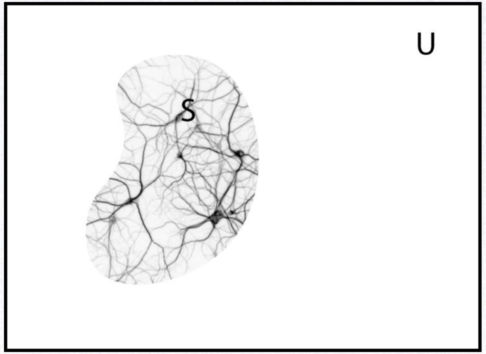

[The Case for Learned Index Structures](#)
[Kraska et al. (2017](#))

DeepMind

Case for Meta-Learning

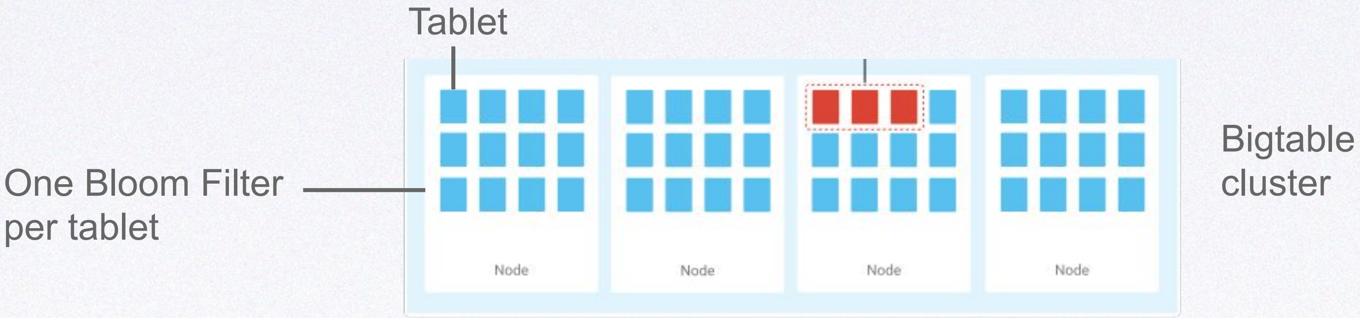

Often data-structures are not created in isolation.

E.g. a Bigtable database with 10,000 tablets.

Common rowkey schema and query distribution.

Meta-learning: slow-learn common distribution, fast-learning of specific set.

Neural Bloom Filter

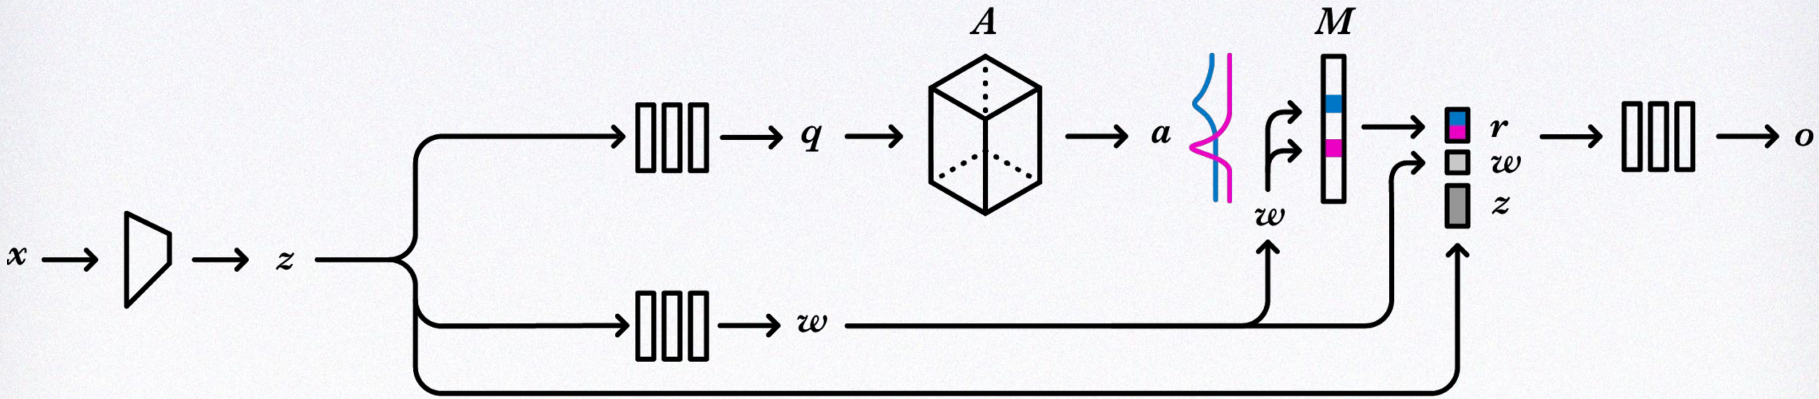

Database Task

	Neural Bloom Filter	Cuckoo Filter
5% FPR	35.8x	0.9x
1% FPR	32.0x	1.2x
0.1% FPR	2.9x	1.2x

Space reduction over Bloom Filter for storage set of 5,000 strings.

Speed Benchmark

	Query + Insert Latency		Query Throughput (QPS)		Insert Throughput (IPS)	
	CPU	GPU	CPU	GPU	CPU	GPU
Bloom Filter [1]	**0.02ms**	-	61K	-	61K	-
Neural Bloom Filter	5.1ms	13ms	3.5K	105K	3.2K	**101K**
LSTM	5.0ms	13ms	3.1K	**107K**	2.4K	4.6K
Learned Index [2]	780ms	1.36s	3.1K	**107K**	25	816

[1] *Query-efficient Bloom Filter* Chen et al. (2007)
[2] *A Case for Learned Index Structures* Kraska et al. (2018)

Talk to me at my poster: <u>**#43**</u>

(Too small to see so you have to come to my poster for the real deal)

More experiments:

Comparisons to MemNets, DNCs, and LSTMs.

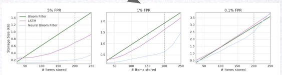

Image tasks with varying structure.

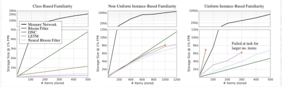

Model ablations to different learned algorithms.