Learning to Reason for Neural Question Answering

Jianfeng Gao

Joint work with **Ming-Wei Chang, Jianshu Chen, Weizhu Chen, Kevin Duh, Yuqing Guo, Po-Sen Huang, Xiaodong Liu, and Yelong Shen.**

Microsoft

MRQA workshop (ACL 2018)

Open-Domain Question Answering (QA)

Will I qualify for OSAP if I'm new in Canada?

Selected Passages from Bing

"Visit the OSAP website for application deadlines. To get OSAP, you have to be eligible. You can apply using an online form, or you can print off the application forms. If you submit a paper application, you must pay an application fee. The online application is free."

Source: http://settlement.org/ontario/education/colleges-universities-and-institutes/financial-assistance-for-post-secondary-education/how-do-i-apply-for-the-ontario-student-assistance-program-osap/

"To be eligible to apply for financial assistance from the Ontario Student Assistance Program (OSAP), you must be a: 1 Canadian citizen; 2 Permanent resident; or 3 Protected person/convention refugee with a Protected Persons Status Document (PPSD)."

Source: http://settlement.org/ontario/education/colleges-universities-and-institutes/financial-assistance-for-post-secondary-education/who-is-eligible-for-the-ontario-student-assistance-program-osap/

"You will not be eligible for a Canada-Ontario Integrated Student Loan, but can apply for a part-time loan through the Canada Student Loans program. There are also grants, bursaries and scholarships available for both full-time and part-time students."

Source: http://www.campusaccess.com/financial-aid/osap.html

Answer
No. You won't qualify.

Text-QA

What is Obama's citizenship?

Selected subgraph from Microsoft's Satori

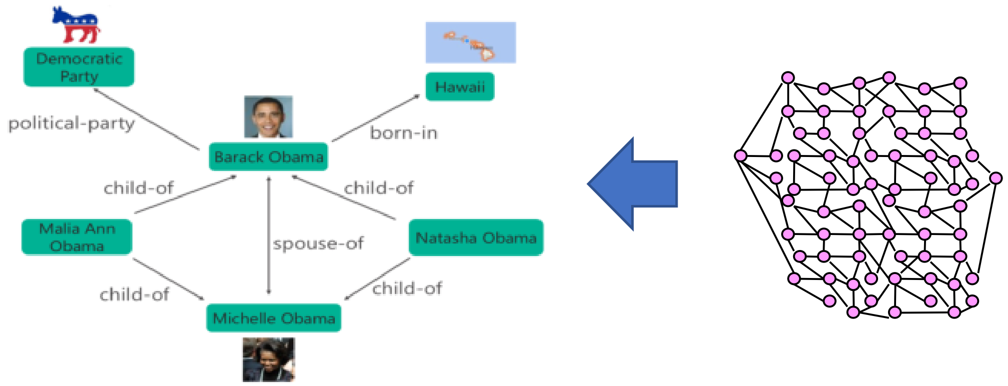

Answer
USA

Knowledge Base (KB)-QA

Question Answering (QA) on Knowledge Base

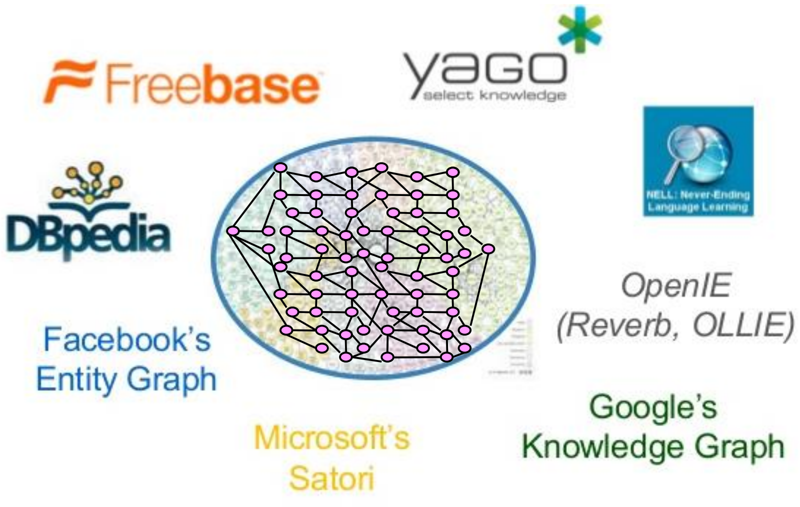

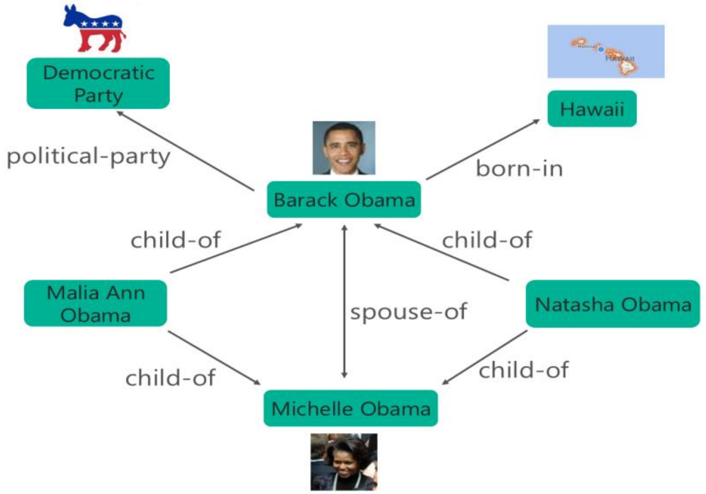

Large-scale knowledge graphs

- Properties of billions of entities
- Plus relations among them

An QA Example:

Question: `what is Obama's citizenship?`

- Query parsing:
  `(Obama, Citizenship,?)`
- Identify and infer over relevant subgraphs:
  `(Obama, BornIn, Hawaii)`
  `(Hawaii, PartOf, USA)`
- correlating semantically relevant relations:
  `BornIn ~ Citizenship`

Answer: `USA`

Reasoning over KG in symbolic vs neural spaces

Symbolic: comprehensible but not robust

- Development: writing/learning production rules
- Runtime : random walk in symbolic space
- E.g., PRA [Lao+ 11], MindNet [Richardson+ 98]

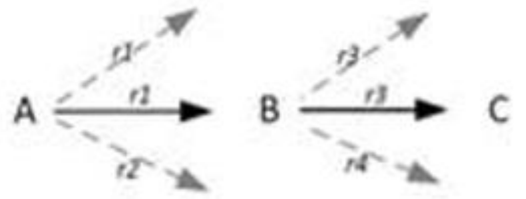

Neural: robust but not comprehensible

- Development: encoding knowledge in neural space
- Runtime : multi-turn querying in neural space (similar to nearest neighbor)
- E.g., ReasoNet [Shen+ 16], DistMult [Yang+ 15]

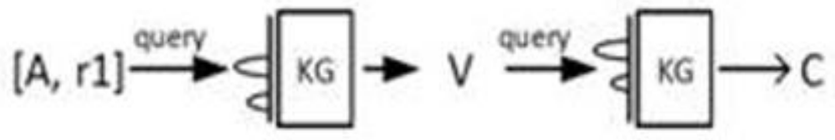

Hybrid: robust and comprehensible

- Development: learning policy π that maps states in neural space to actions in symbolic space via RL
- Runtime : graph walk in symbolic space guided by π
- E.g., M-Walk [Shen+ 18], DeepPath [Xiong+ 18], MINERVA [Das+ 18]

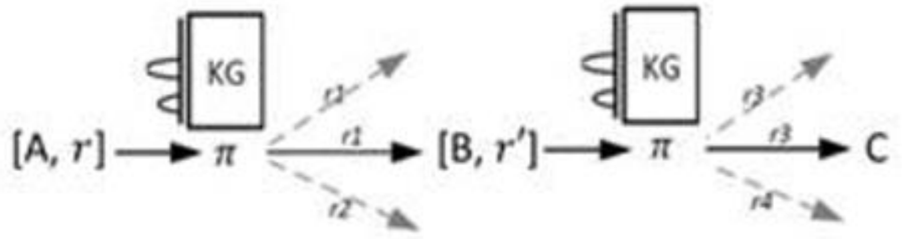

Symbolic approaches to QA

- Understand the question via **semantic parsing**
  - Input: what is Obama's citizenship?
  - Output (LF): (Obama, Citizenship,?)
- Collect relevant information via fuzzy **keyword matching**
  - (Obama, BornIn, Hawaii)
  - (Hawaii, PartOf, USA)
  - Needs to know that BornIn and Citizenship are semantically related
- Generate the answer via **reasoning**
  - (Obama, Citizenship, **USA**)
- **Challenges**
  - **Paraphrasing in NL**
  - **Search complexity of a big KG**

[Richardson+ 98; Berant+ 13; Yao+ 15; Bao+ 14; Yih+ 15; etc.]

Key Challenge in KB-QA: *Language Mismatch (Paraphrasing)*

- Lots of ways to ask the same question
  - *“What was the date that Minnesota became a state?”*
  - *“Minnesota became a state on?”*
  - *“When was the state Minnesota created?”*
  - *“Minnesota's date it entered the union?”*
  - *“When was Minnesota established as a state?”*
  - *“What day did Minnesota officially become a state?”*

- Need to map them to the predicate defined in KB
  - location.dated_location.date_founded

Scaling up semantic parsers

- Paraphrasing in NL
  - Introduce a paragraphing engine as pre-processor [Berant&Liang 14]
  - Using semantic similarity model (e.g., DSSM) for semantic matching [Yih+ 15]
- Search complexity of a big KG
  - Pruning (partial) paths using domain knowledge
- More details: IJCAI-2016 tutorial on “Deep Learning and Continuous Representations for Natural Language Processing” by Yih, He and Gao.

From symbolic to neural computation

Input: Q

Symbolic → Neural
by **Encoding** (Q/D/Knowledge)

Symbolic Space
- human readable

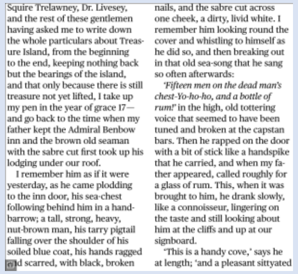
Squire Trelawney, Dr. Livesey, and the rest of these gentlemen having asked me to write down the whole particulars about Treasure Island, from the beginning to the end, keeping nothing back but the bearings of the island, and that only because there is still treasure not yet lifted, I take up my pen in the year of grace 17— and go back to the time when my father kept the Admiral Benbow inn and the brown old seaman with the sabre cut first took up his lodging under our roof.

I remember him as if it were yesterday, as he came plodding to the inn door, his sea-chest following behind him in a hand-barrow; a tall, strong, heavy, nut-brown man, his tarry pigtail falling over the shoulder of his soiled blue coat, his hands ragged and scarred, with black, broken nails, and the sabre cut across one cheek, a dirty, livid white. I remember him looking round the cover and whistling to himself as he did so, and then breaking out in that old sea-song that he sang so often afterwards:

'Fifteen men on the dead man's chest-Yo-ho-ho, and a bottle of rum!' in the high, old tottering voice that seemed to have been tuned and broken at the capstan bars. Then he rapped on the door with a bit of stick like a handspike that he carried, and when my father appeared, called roughly for a glass of rum. This, when it was brought to him, he drank slowly, like a connoisseur, lingering on the taste and still looking about him at the cliffs and up at our signboard.

'This is a handy cove,' says he at length; 'and a pleasant sittyated

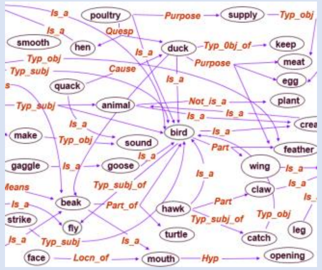

Reasoning: Question + KB → answer vector via multi-step inference, summarization, deduction etc.

Neural Space
- Computationally efficient

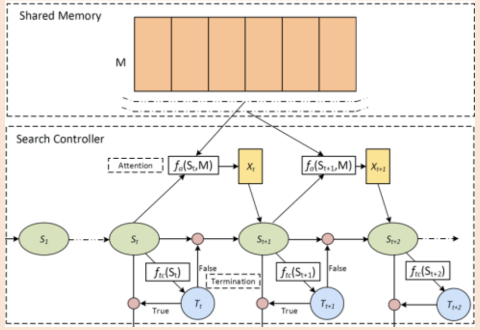

Error(A, A*)

Output: A

Neural → Symbolic
by **Decoding** (synthesizing answer)

Case study: ReasoNet with Shared Memory

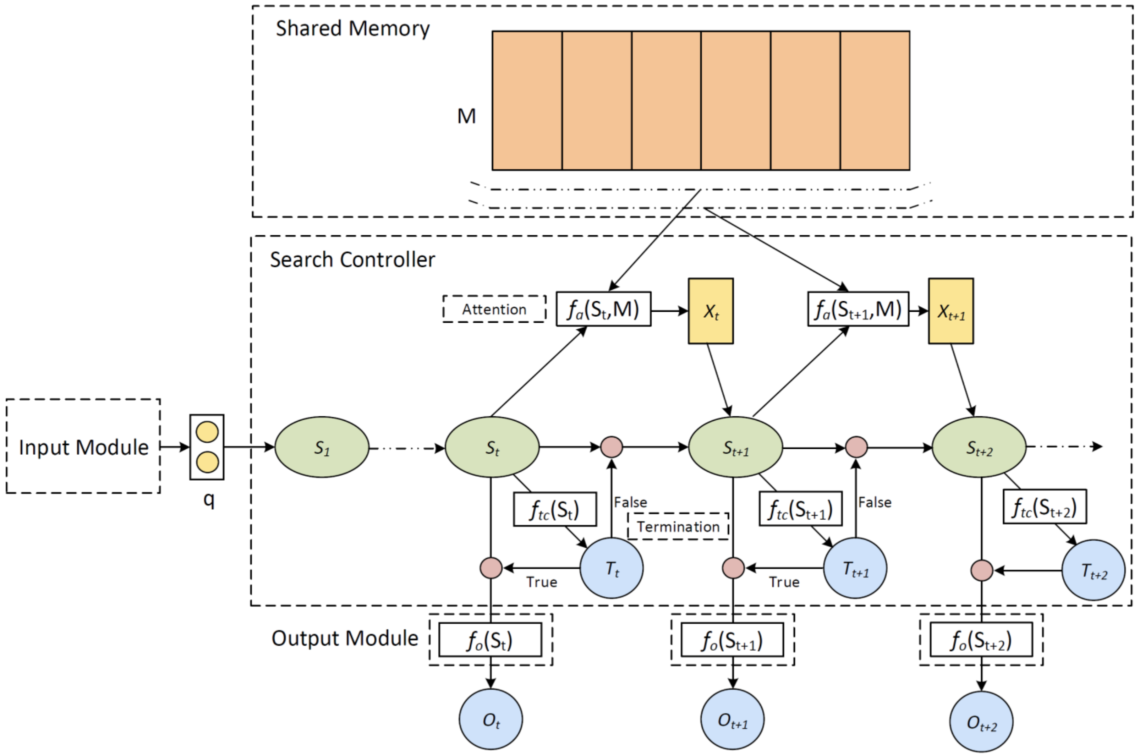

- **Shared memory** (M) encodes task-specific knowledge
  - **Long-term memory:** encode KB for answering all questions in QA on KB
  - **Short-term memory:** encode the passage(s) which contains the answer of a question in QA on Text
- **Working memory** (hidden state S_t) contains a description of the current state of the world in a reasoning process
- **Search controller** performs multi-step inference to update S_t of a question using knowledge in shared memory
- Input/output modules are task-specific

[Shen+ 16; Shen+ 17]

Joint learning of Shared Memory and Search Controller

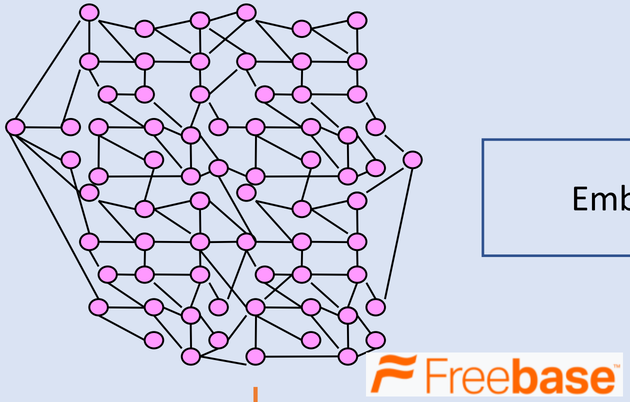

Embed KG to memory vectors

Paths extracted from KG:

(John, BornIn, Hawaii)

(Hawaii, PartOf, USA)

(John, *Citizenship*, USA)

…

Training samples generated

(John, BornIn, ?)->(Hawaii)

(Hawaii, PartOf, ?)->(USA)

(John, Citizenship, ?)->(USA)

…

(John, Citizenship, ?)

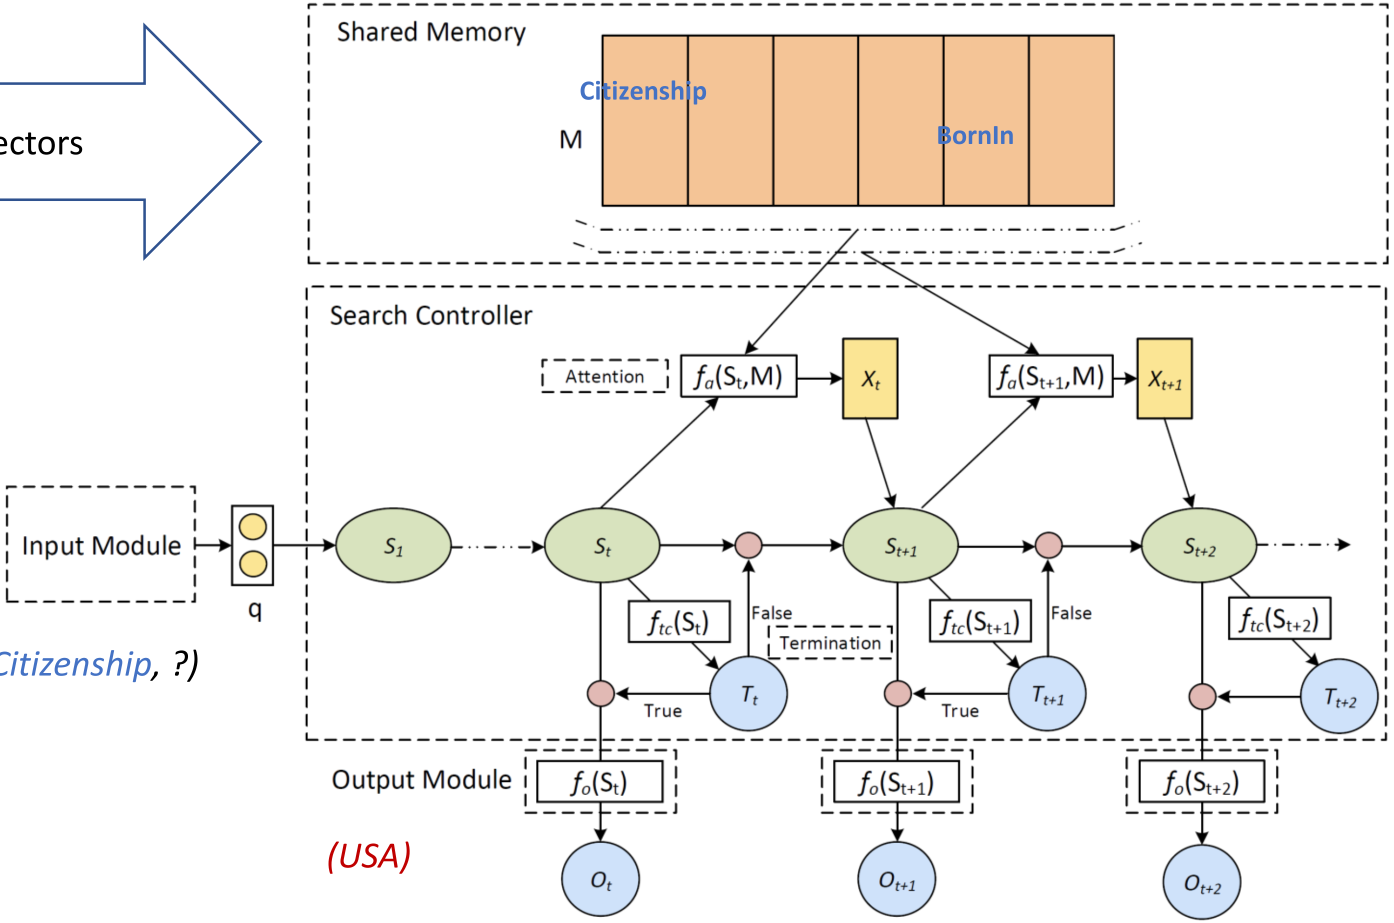

Joint learning of Shared Memory and Search Controller

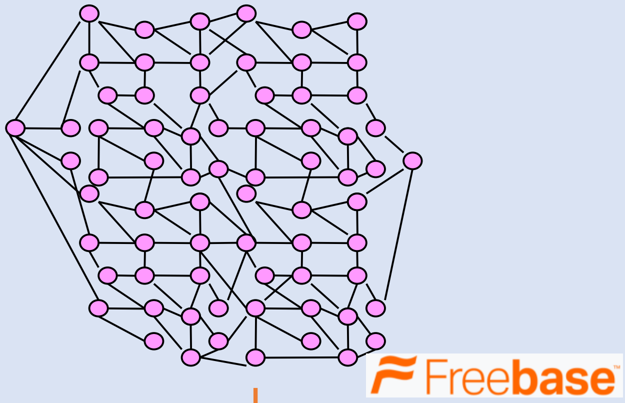

Paths extracted from KG:

(John, BornIn, Hawaii)
(Hawaii, PartOf, USA)
(John, *Citizenship*, USA)
…

Training samples generated

(John, BornIn, ?)->(Hawaii)
(Hawaii, PartOf, ?)->(USA)
(John, Citizenship, ?)->(USA)
…

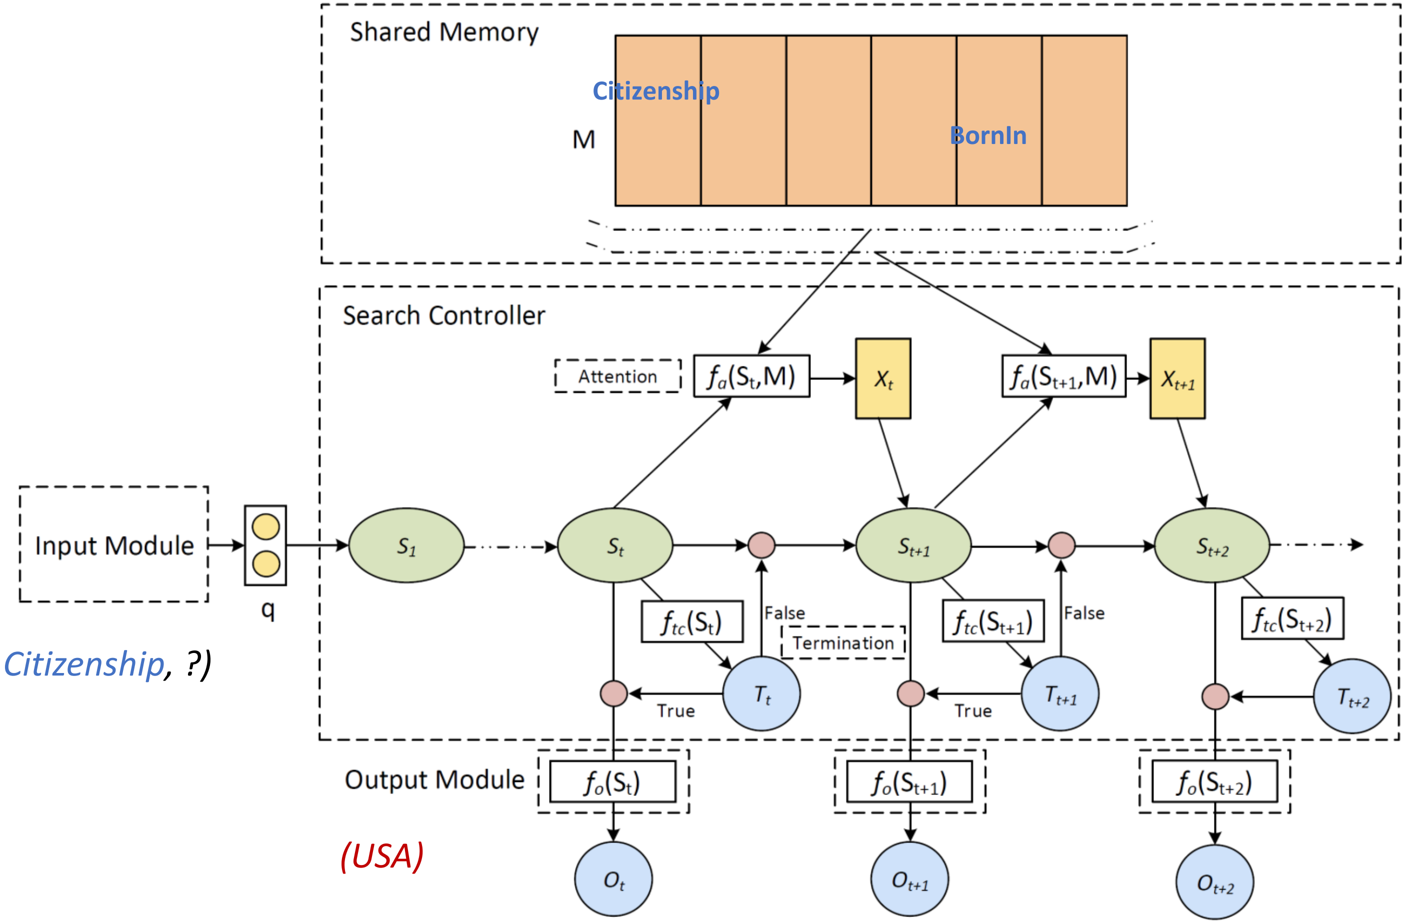

(John, Citizenship, ?)

(USA)

Shared Memory: long-term memory to store learned knowledge, like human brain

- Knowledge is learned via performing tasks, e.g., update memory to answer new questions
- New knowledge is *implicitly* stored in memory cells via gradient update
- Semantically relevant relations/entities can be compactly represented using similar vectors.

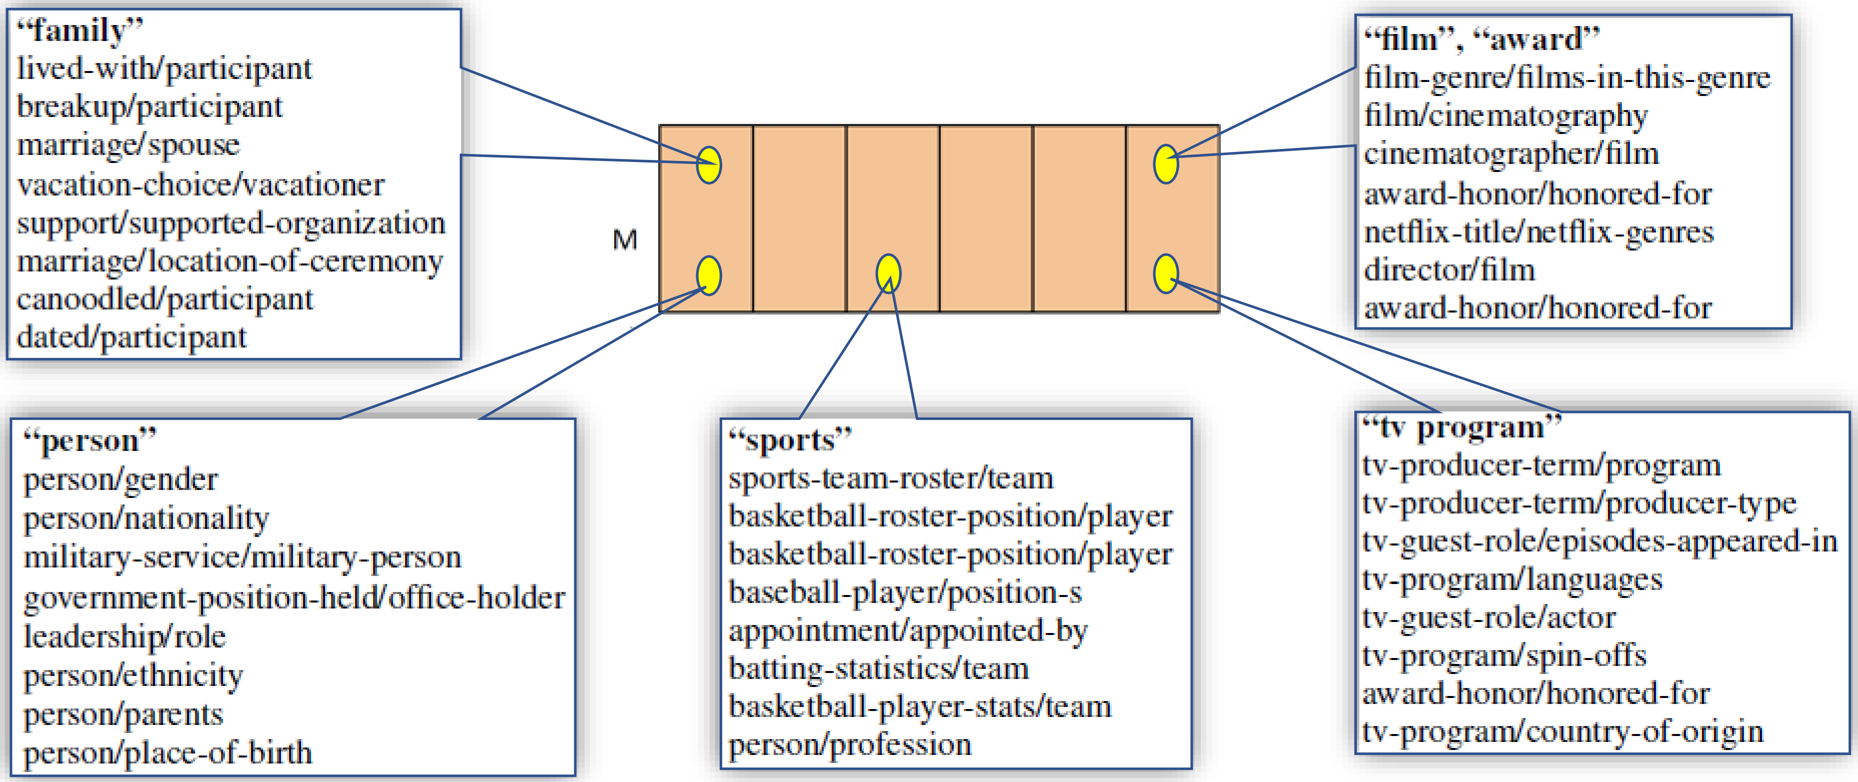

Search controller for KB QA

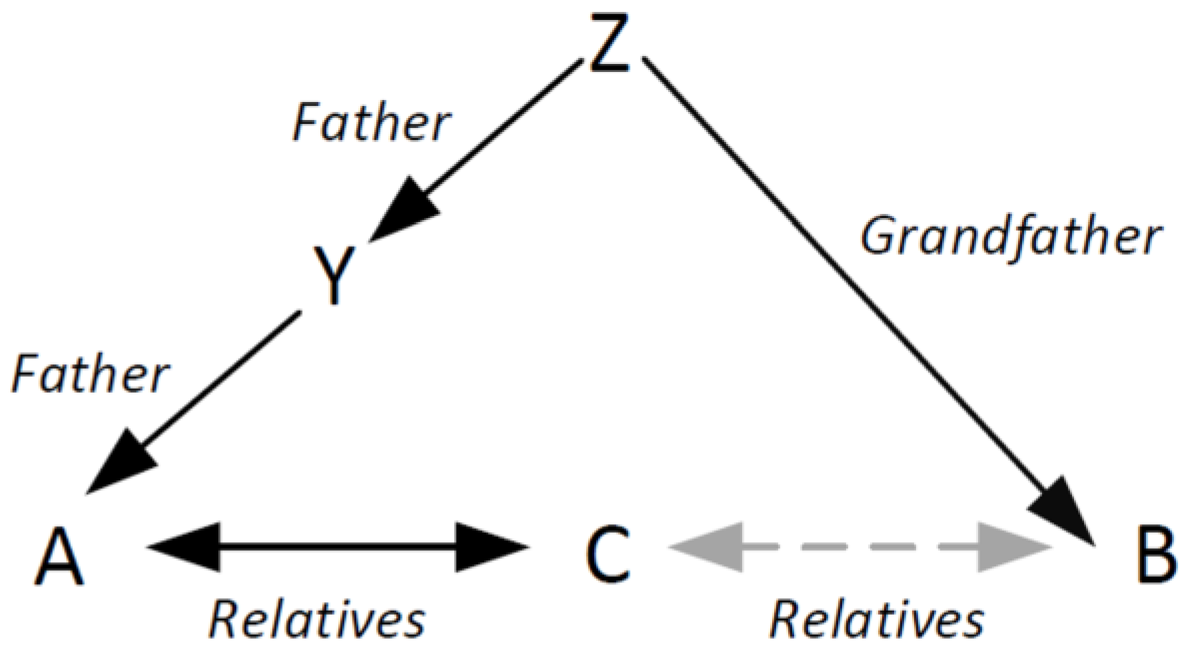

Space	Inference Path
Symbolic	$B \xrightarrow{\text{Grandfather}^{-1}} Z \xrightarrow{\text{Father}} Y \xrightarrow{\text{Father}} A \xrightarrow{\text{Relatives}} C$
Neural	$B \xrightarrow{\text{Relatives}} ? \Rightarrow A \xrightarrow{\text{Relatives}} C$

[Shen+ 16]

M-Walk: Learning to Reason over Knowledge Graph

- Graph Walking as a Markov Decision Process
  - State: encode "traversed nodes + previous actions + initial query" using RNN
  - Action: choose an edge and move to the next node, or STOP
  - Reward: +1 if stop at a correct node, 0 otherwise
  - Learning to reason over KG = seeking an optimal policy π

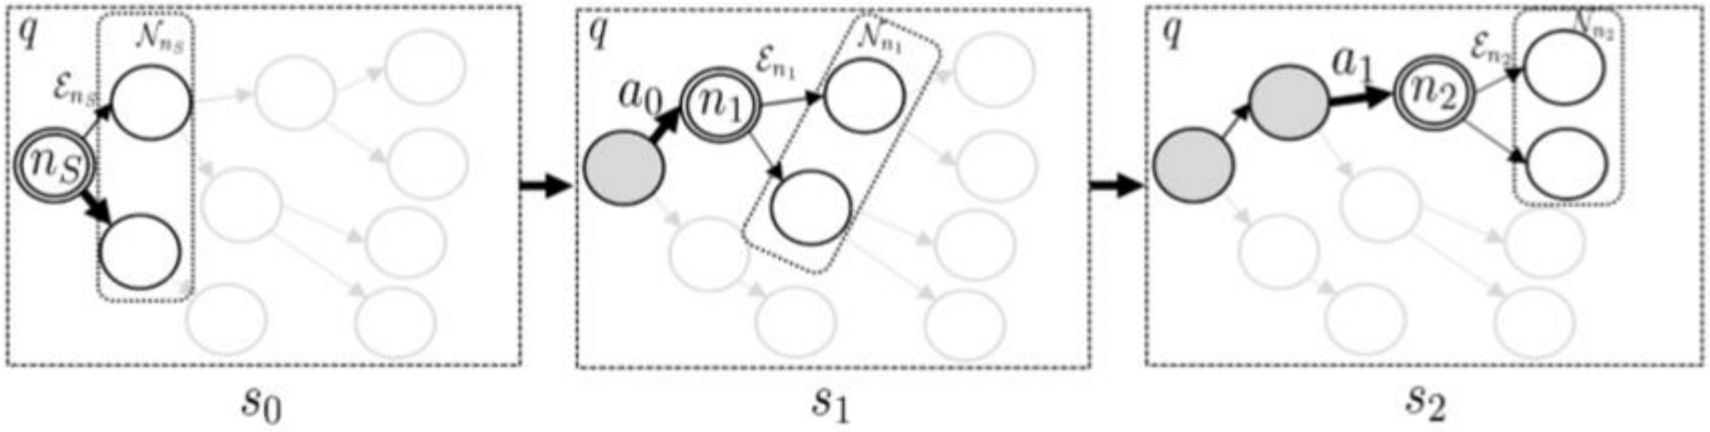

Training with Monte Carlo Tree Search (MCTS)

- Address sparse reward by *running MCTS simulations to generate trajectories with more positive reward*
- Exploit that KG is given and MDP transitions are deterministic
- On each MCTS simulation, roll out a trajectory by selecting actions
  - Treat π as a prior
  - Prefer actions with high value (i.e., $\frac{W(s,a)}{N(s,a)}$, where N and W are visit count and action reward estimated using value network)

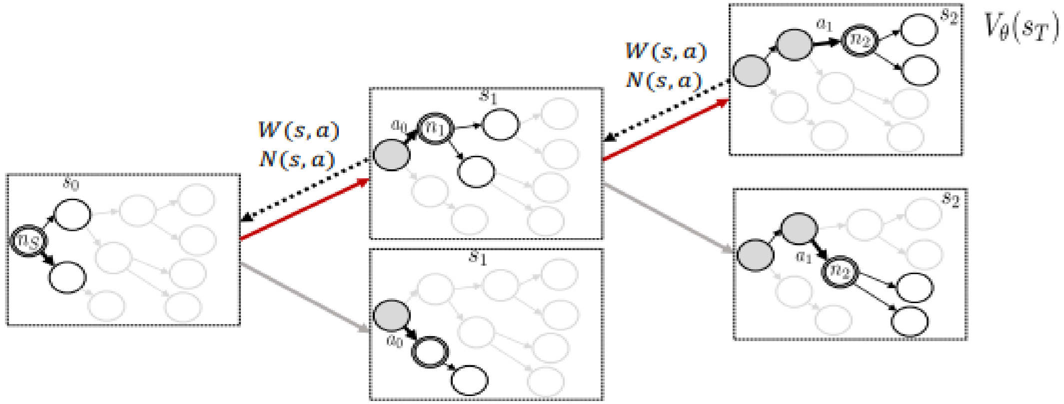

(a) An example of MCTS path (in red) in M-Walk

Joint learning of π_θ, V_θ, and Q_θ

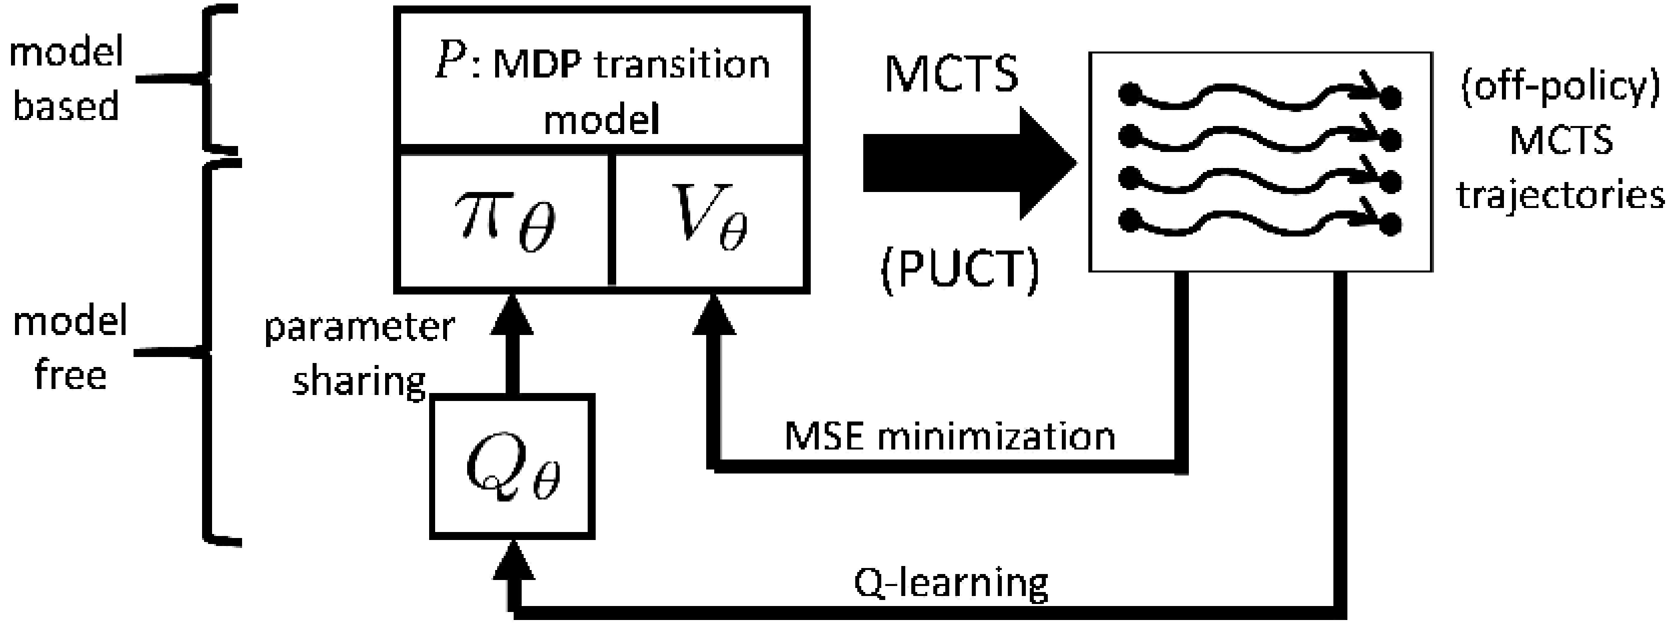

(b) Iterative policy improvement in M-Walk

Experiments on NELL-995

- NELL-995 dataset:
  - 154,213 Triples
  - 75,492 unique entities
  - 200 unique relations.
- Missing link prediction Task:
  - Predict the tail entity given the head entity and relation
  - i.e., Citizenship (Obama, ?) → USA
- Evaluation Metric:
  - Mean Average Precision (the higher the better)

Missing Link Prediction Results

Table 1. NELL-995 Link Prediction Performance Comparison using MAP scores.

Tasks	PRA
athletePlaysForTeam	0.547
athletePlaysInLeague	0.841
athleteHomeStadium	0.859
athletePlaysSport	0.474
teamPlaySports	0.791
orgHeadquaterCity	0.811
worksFor	0.681
bornLocation	0.668
personLeadsOrg	0.700
orgHiredPerson	0.599
Overall	0.697

Path Ranking Algorithm:
Symbolic Reasoning Approach

Missing Link Prediction Results

Table 1. NELL-995 Link Prediction Performance Comparison using MAP scores.

Tasks	PRA	TransE	TransR	ReasoNet
athletePlaysForTeam	0.547	0.627	0.673	0.789
athletePlaysInLeague	0.841	0.773	0.912	0.936
athleteHomeStadium	0.859	0.718	0.722	0.787
athletePlaysSport	0.474	0.876	0.963	0.969
teamPlaySports	0.791	0.761	0.814	0.833
orgHeadquaterCity	0.811	0.620	0.657	0.835
worksFor	0.681	0.677	0.692	0.769
bornLocation	0.668	0.712	0.812	**0.836**
personLeadsOrg	0.700	0.751	0.772	0.802
orgHiredPerson	0.599	0.719	0.737	0.768
Overall	0.697	0.723	0.775	0.817

Neural Reasoning Approaches

Path Ranking Algorithm:
Symbolic Reasoning Approach

Missing Link Prediction Results

Table 1. NELL-995 Link Prediction Performance Comparison using MAP scores.

Tasks	ReinforceWalk	PG	A2C	MINERVA[3]	DeepPath	PRA	TransE	TransR	ReasoNet
athletePlaysForTeam	**0.831**	0.769	0.700	0.630	0.750	0.547	0.627	0.673	0.789
athletePlaysInLeague	**0.974**	0.955	0.955	0.837	0.960	0.841	0.773	0.912	0.936
athleteHomeStadium	**0.905**	0.865	0.861	0.557	0.890	0.859	0.718	0.722	0.787
athletePlaysSport	**0.985**	0.962	0.971	0.916	0.957	0.474	0.876	0.963	0.969
teamPlaySports	**0.881**	0.631	0.679	0.751	0.738	0.791	0.761	0.814	0.833
orgHeadquaterCity	0.943	0.935	0.928	0.947	0.790	0.811	0.620	0.657	0.835
worksFor	**0.786**	0.758	0.758	0.752	0.711	0.681	0.677	0.692	0.769
bornLocation	0.786	0.767	0.766	0.782	0.757	0.668	0.712	0.812	**0.836**
personLeadsOrg	**0.821**	0.802	0.810	0.771	0.795	0.700	0.751	0.772	0.802
orgHiredPerson	0.843	0.832	0.839	0.860	0.742	0.599	0.719	0.737	0.768
Overall	**0.876**	0.828	0.827	0.780	0.809	0.697	0.723	0.775	0.817

Two variants of ReinforceWalk without MCTS

Reinforcement Symbolic + Neural Reasoning Approaches

Path Ranking Algorithm:
Symbolic Reasoning Approach

Neural Reasoning Approaches

Neural MRC Models on SQuAD

What types of European groups were able to avoid the plague?

From Italy, the disease spread northwest across Europe, striking France, Spain, Portugal and England by June 1348, then turned and spread east through Germany and Scandinavia from 1348 to 1350. It was introduced in Norway in 1349 when a ship landed at Askøy, then spread to Bjørgvin (modern Bergen) and Iceland. Finally it spread to northwestern Russia in 1351. The plague was somewhat less common in parts of Europe that had smaller trade relations with their neighbours, including the Kingdom of Poland, the majority of the Basque Country, isolated parts of Belgium and the Netherlands, and isolated alpine villages throughout the continent.

A limited form of comprehension:

- No need for extra knowledge outside the paragraph
- No need for clarifying questions
- The answer must exist in the paragraph
- The answer must be a text span, not synthesized

- Encoding: map each text span to a semantic vector
- *Reasoning: rank and re-rank semantic vectors*
- *Decoding: map the top-ranked vector to text*

Neural MRC models…

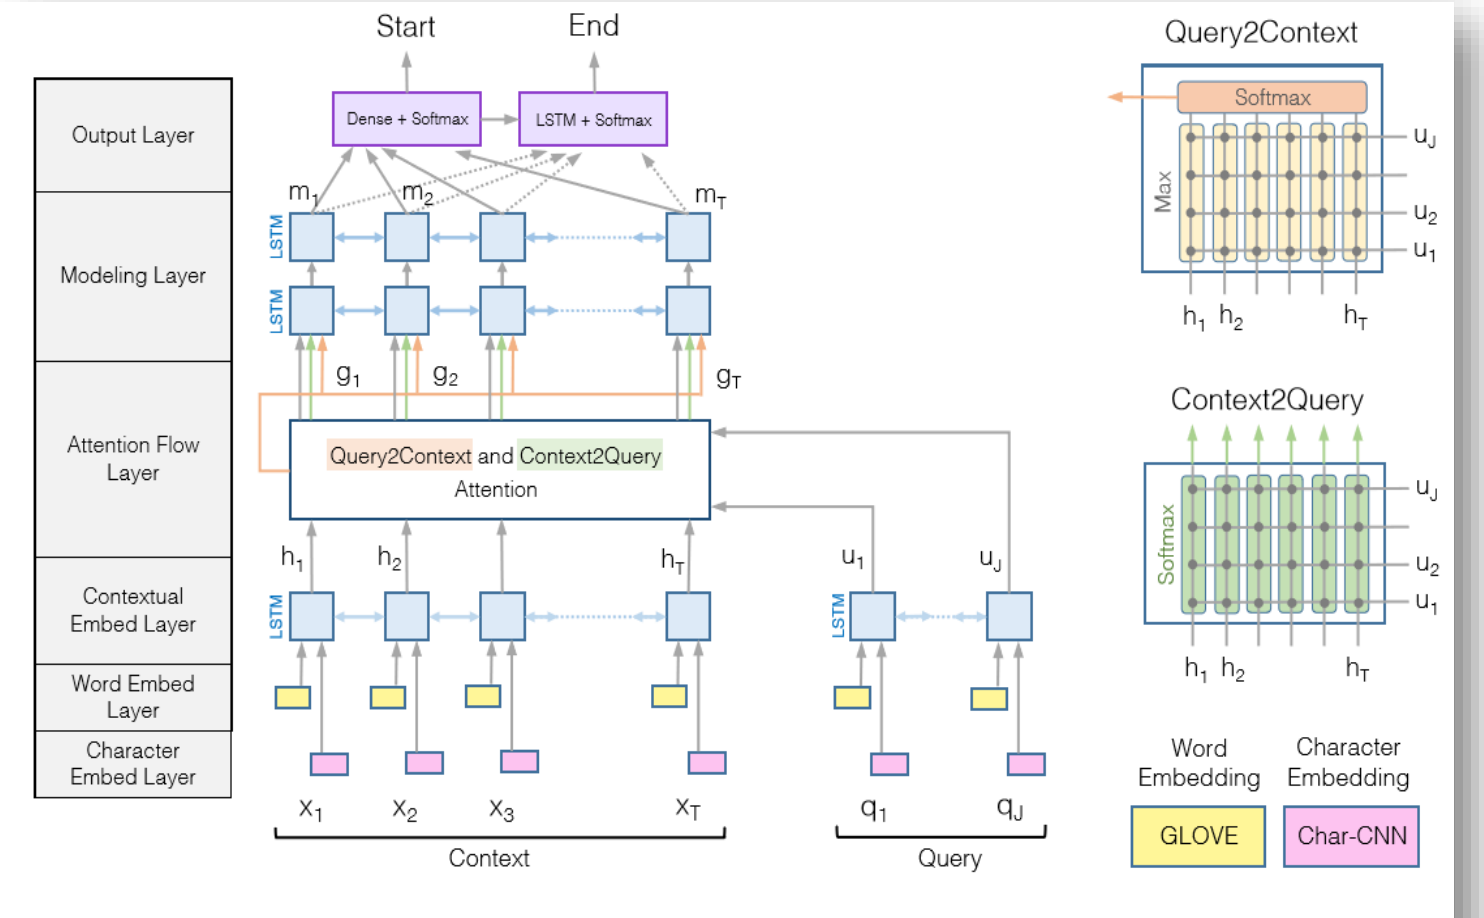

Figure 1: BiDirectional Attention Flow Model *(best viewed in color)*

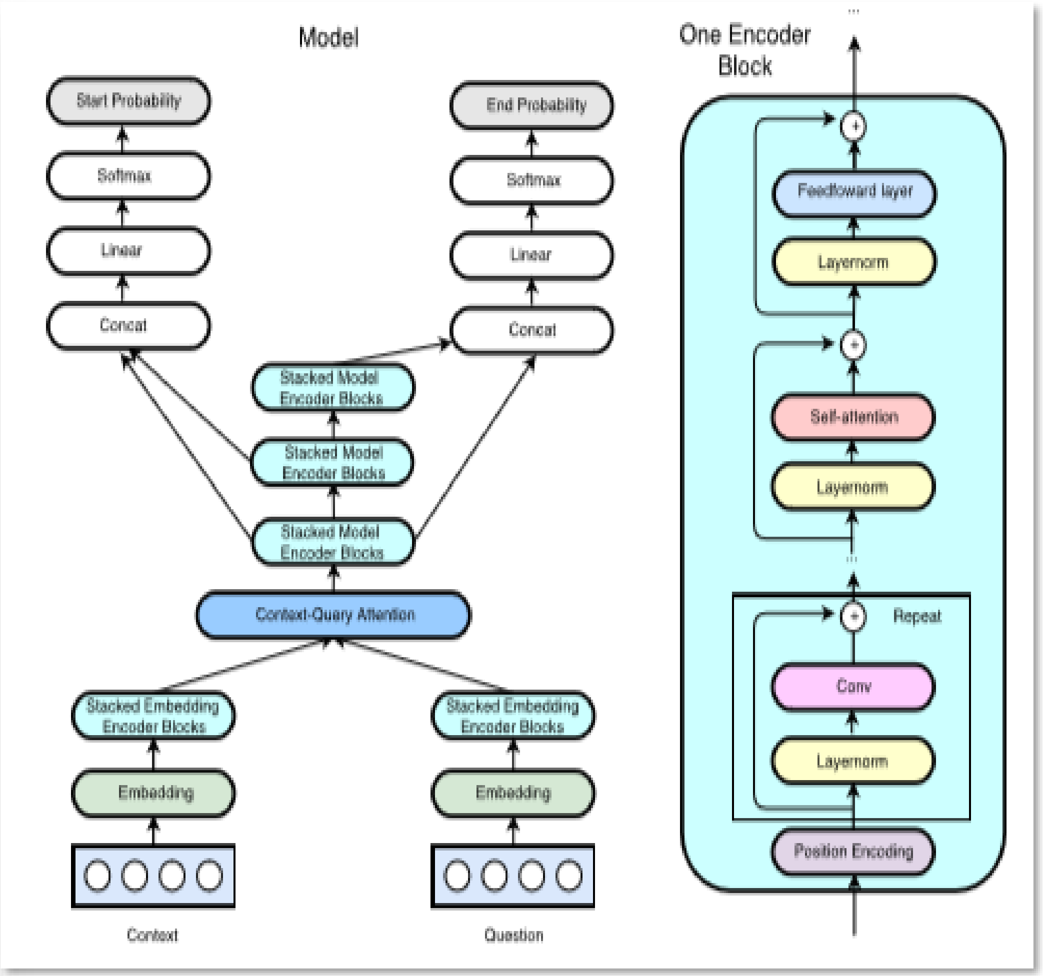

[Seo+ 16; Yu+ 18]

Text-QA

Will I qualify for OSAP if I'm new in Canada?

Selected Passages from Bing

"Visit the OSAP website for application deadlines. To get OSAP, you have to be eligible. You can apply using an online form, or you can print off the application forms. If you submit a paper application, you must pay an application fee. The online application is free."

Source: http://settlement.org/ontario/education/colleges-universities-and-institutes/financial-assistance-for-post-secondary-education/how-do-i-apply-for-the-ontario-student-assistance-program-osap/

"To be eligible to apply for financial assistance from the Ontario Student Assistance Program (OSAP), you must be a: 1 Canadian citizen; 2 Permanent resident; or 3 Protected person/convention refugee with a Protected Persons Status Document (PPSD)."

Source: http://settlement.org/ontario/education/colleges-universities-and-institutes/financial-assistance-for-post-secondary-education/who-is-eligible-for-the-ontario-student-assistance-program-osap/

"You will not be eligible for a Canada-Ontario Integrated Student Loan, but can apply for a part-time loan through the Canada Student Loans program. There are also grants, bursaries and scholarships available for both full-time and part-time students."

Source: http://www.campusaccess.com/financial-aid/osap.html

Answer

No. You won't qualify.

MS MARCO [Nguyen+ 16]

In meteorology, precipitation is any product of the condensation of atmospheric water vapor that falls under gravity. The main forms of precipitation include drizzle, rain, sleet, snow, graupel and hail... Precipitation forms as smaller droplets coalesce via collision with other rain drops or ice crystals within a cloud. Short, intense periods of rain in scattered locations are called "showers".

What causes precipitation to fall?
gravity

What is another main form of precipitation besides drizzle, rain, snow, sleet and hail?
graupel

Where do water droplets collide with ice crystals to form precipitation?
within a cloud

Figure 1: Question-answer pairs for a sample passage in the SQuAD dataset. Each of the answers is a segment of text from the passage.

SQuAD [Rajpurkar+ 16]

MS MARCO

Microsoft MAchine Reading COmprehension Dataset

Follow MSMarcoAI

MSMARCO V2.1 is now availible. Start Experimenting today!

1,010,916 Real Bing User Queries

182,669 Natural Language Answers

No Answer Subset

10 Passages Per Query

Sign Up For Email Updates | Read Our Paper | Join Us On Slack

Microsoft

Novice Task

Rank	Model	Submission Date	Rouge-L	Bleu-1	F1
1	**Human Performance**	April 23th, 2018	53.87	48.50	94.72
2	**VNET** Baidu NLP	June 19th, 2018	46.72	50.45	70.96
3	**SNET** JY Zhao	June 26th, 2018	42.36	46.14	70.96
4	**DNET++** QA Geeks	June 1st, 2018	41.91	45.80	70.93
5	**SNET+seq2seq** Yihan Ni of the CAS Key Lab of Web Data Science and Technology, ICT, CAS	June 1st, 2018	39.82	42.27	70.96
7	**DNET** QA Geeks	May 29th, 2018	33.30	29.12	74.36
8	**BIDAF+seq2seq** Yihan Ni of the CAS Key Lab of Web Data Science and Technology, ICT, CAS	May 29th, 2018	27.60	28.84	70.96
9	**BiDaF Baseline(Implemented By MSMARCO Team)** Allen Institute for AI & University of Washington [Seo et al. '16]	April 23th, 2018	23.96	10.64	74.93

Intermediate Task

Rank	Model	Submission Date	Rouge-L	Bleu-1
1	**Human Performance**	April 23th, 2018	63.21	53.03
2	**VNET** Baidu NLP	July 4th, 2018	46.41	43.12
3	**ConZNet** S3R	July 16th, 2018	42.14	38.62
4	**Bayes QA** Bin Bi of Alibaba NLP	June 14st, 2018	41.11	43.54
5	**SNET+seq2seq** Yihan Ni of the CAS Key Lab of Web Data Science and Technology, ICT, CAS	June 1st, 2018	40.07	37.54
6	**BIDAF+seq2seq** Yihan Ni of the CAS Key Lab of Web Data Science and Technology, ICT, CAS	May 29th, 2018	32.22	28.33
7	**DNET++** QA Geeks	June 1st, 2018	26.15	32.24
8	**DNET** QA Geeks	May 29th, 2018	25.19	30.73
9	**SNET** JY Zhao	May 29th, 2018	24.66	30.78
10	**BiDaF Baseline(Implemented By MSMARCO Team)** Allen Institute for AI & University of Washington [Seo et al. '16]	April 23th, 2018	16.91	9.30

Multi-step reasoning: example

Query	Who was the #2 pick in the 2011 NFL Draft?
Passage	Manning was the #1 selection of the 1998 NFL draft, while Newton was picked first in 2011. The matchup also pits the top two picks of the 2011 draft against each other: Newton for Carolina and **Von Miller** for Denver.
Answer	**Von Miller**

- Step 1:
  - **Extract:** Manning is #1 pick of 1998
  - **Infer:** Manning is NOT the answer
- Step 2:
  - **Extract:** Newton is #1 pick of 2011
  - **Infer:** Newton is NOT the answer
- Step 3:
  - **Extract:** Newton and Von Miller are top 2 picks of 2011
  - **Infer:** **Von Miller** is the #2 pick of 2011

ReasoNet: learn to stop reading

With Q in mind, read Doc repeatedly, each time focusing on different parts of doc until a satisfied answer is formed:

1. Given a set of docs in memory: $\mathbf{M}$
2. Start with query: S
3. Identify info in $\mathbf{M}$ that is related to S : $X = f_a(S, \mathbf{M})$
4. Update internal state: $S = \mathrm{RNN}(S, X)$
5. Whether a satisfied answer O can be formed based on S: $f_{tc}(S)$
6. If so, stop and output answer $O = f_o(S)$; otherwise return to 3.

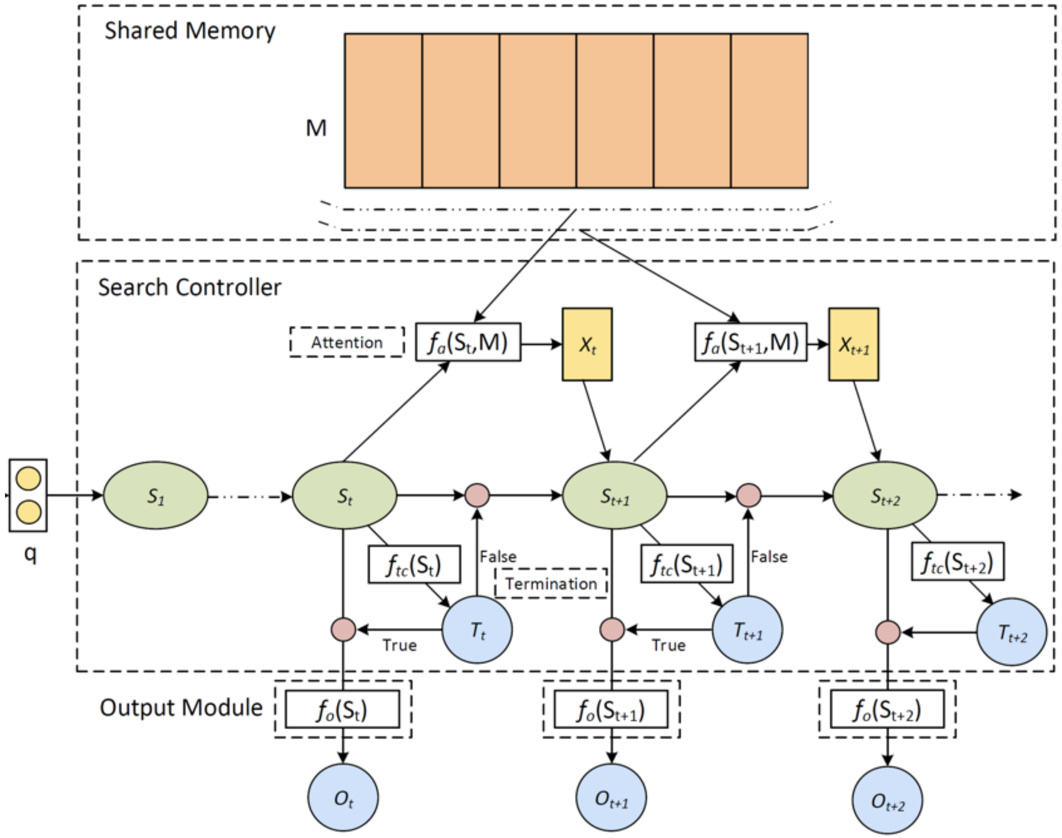

The step size is determined dynamically based on the complexity of the problem using reinforcement learning.

[Shen+ 17]

ReasoNet: learn to stop reading

Query	Who was the #2 pick in the 2011 NFL Draft?
Passage	Manning was the #1 selection of the 1998 NFL draft, while Newton was picked first in 2011. The matchup also pits the top two picks of the 2011 draft against each other: Newton for Carolina and Von Miller for Denver.
Answer	**Von Miller**

Rank-1
Rank-2
Rank-3

S: Who was the #2 pick in the 2011 NFL Draft?

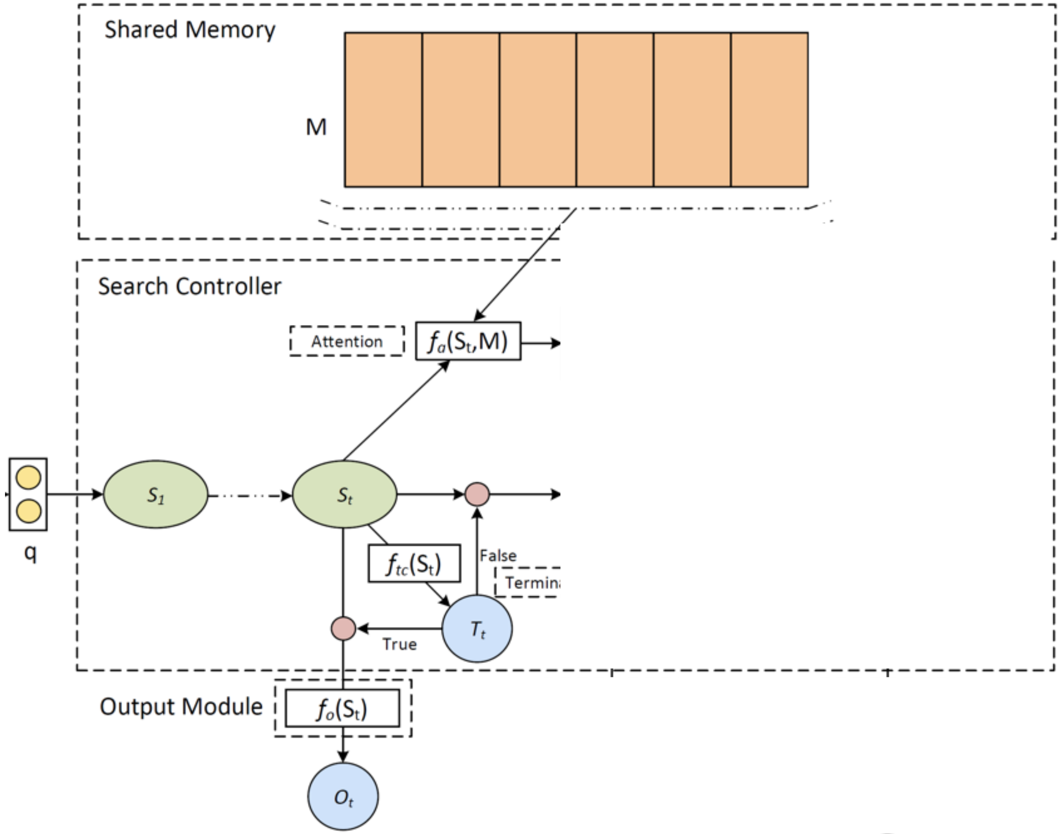

Step	Termination Probability	Prob. Answer
1	0.001	0.392

ReasoNet: learn to stop reading

Query	Who was the #2 pick in the 2011 NFL Draft?
Passage	Manning was the #1 selection of the 1998 NFL draft, while Newton was picked first in 2011. The matchup also pits the top two picks of the 2011 draft against each other: Newton for Carolina and Von Miller for Denver.
Answer	**Von Miller**

Rank-1
Rank-2
Rank-3

S: Manning is #1 pick of 1998, but this is unlikely the answer.

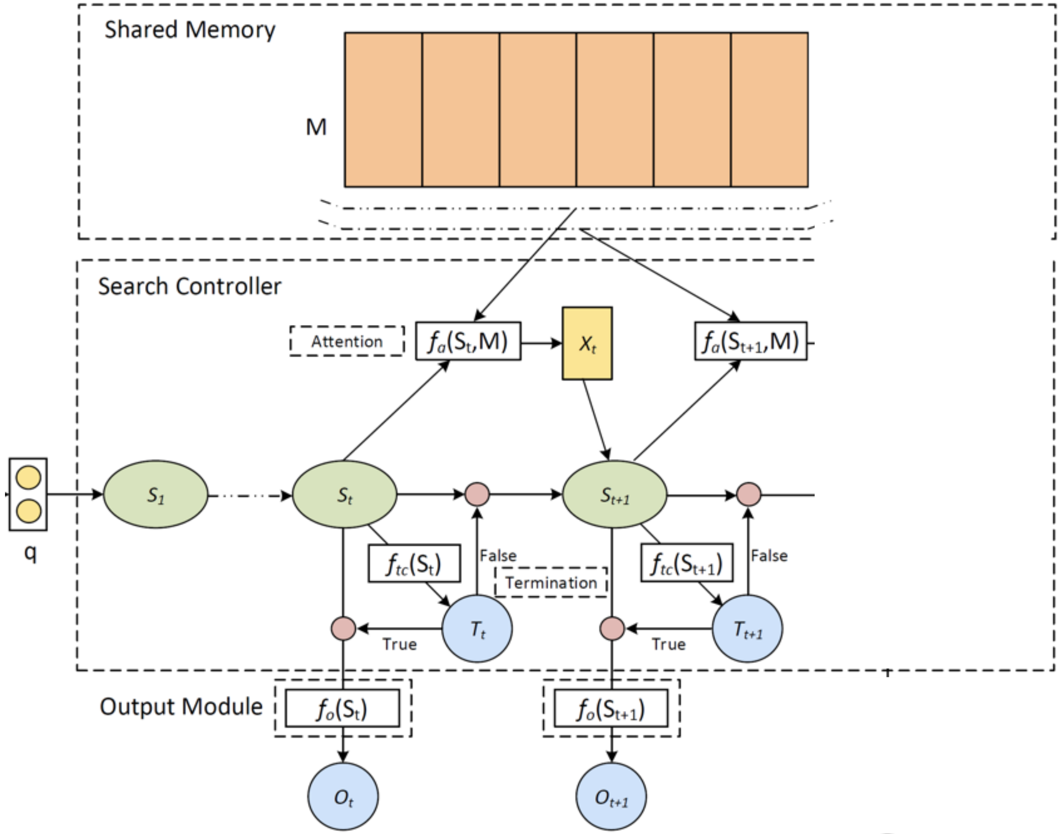

Step	Termination Probability	Prob. Answer
1	0.001	0.392
2	0.675	0.649

ReasoNet: learn to stop reading

Query	Who was the #2 pick in the 2011 NFL Draft?
Passage	Manning was the #1 selection of the 1998 NFL draft, while Newton was picked first in 2011. The matchup also pits the top two picks of the 2011 draft against each other: Newton for Carolina and Von Miller for Denver.
Answer	**Von Miller**

Rank-1
Rank-2
Rank-3

S: Manning is #1 pick of 1998, Newton is #1 pick of 2011, but neither is the answer.

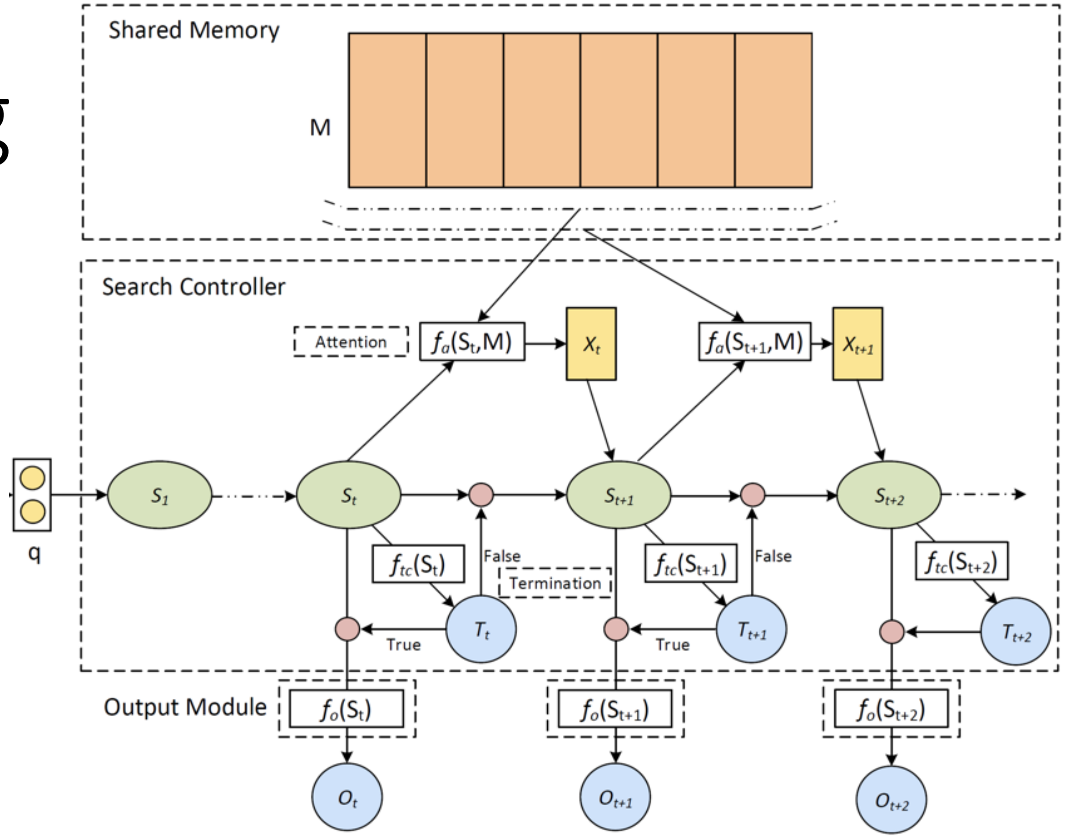

Step t	Termination Probability f_{tc}	Prob. Answer f_o
1	0.001	0.392
2	0.675	0.649
3	0.939	0.865

Stochastic Answer Net

- Training uses *stochastic prediction dropout* on the answer module
- Reasoning employs all the outputs of multiple-step reasoning via voting
- Differs from ReasoNet
  - Easy to train, BP vs. policy gradient
  - Better performance, i.e., best documented MRC model on the SQuAD leaderboard as of Dec. 19, 2017

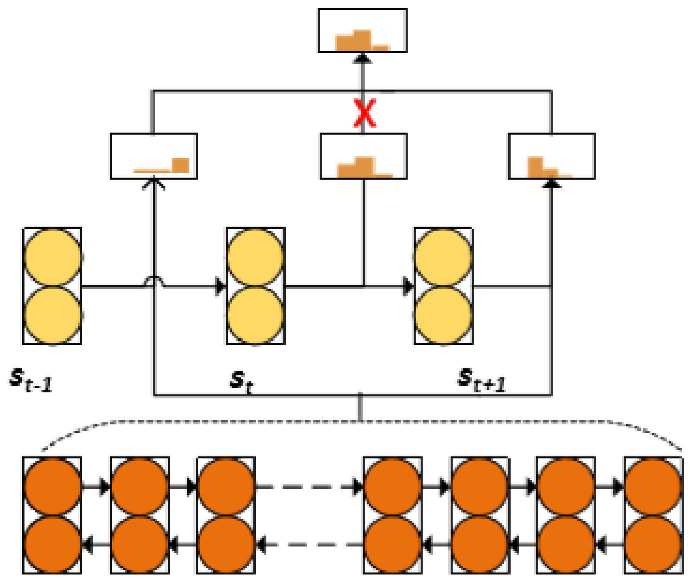

Figure 1: Illustration of "stochastic prediction dropout" in the answer module during training. At each reasoning step t, the model combines memory (bottom row) with hidden states s_{t-1} to generate a prediction (multinomial distribution). Here, there are three steps and three predictions, but one prediction is dropped and the final result is an average of the remaining distributions.

[Liu+ 18]

Answer Module	EM	F1
Standard 1-step	75.139	83.367
Fixed 5-step with Memory Network (prediction from final step)	75.033	83.327
Fixed 5-step with Memory Network (prediction averaged from all steps)	75.256	83.215
Dynamic steps (max 5) with ReasoNet	75.355	83.360
Stochastic Answer Network (SAN), Fixed 5-step	**76.235**	**84.056**

Table 1: SQuAD devset results

SingleModel	ROUGE	BLEU
ReasoNet++(Shen et al., 2017)	38.01	38.62
V-Net(Wang et al., 2018)	45.65	-
Standard 1-step in Table 1	42.30	42.39
SAN	**46.14**	**43.85**

Table 7: **MS MARCO devset results.**

Single model:	AddSent	AddOneSent
LR (Rajpurkar et al., 2016)	23.2	30.3
SEDT (Liu et al., 2017a)	33.9	44.8
BiDAF (Seo et al., 2016)	34.3	45.7
jNet (Zhang et al., 2017)	37.9	47.0
ReasoNet(Shen et al., 2017)	39.4	50.3
RaSoR(Lee et al., 2016)	39.5	49.5
Mnemonic(Hu et al., 2017)	**46.6**	56.0
QANet(Yu et al., 2018)	45.2	55.7
Standard 1-step in Table 1	45.4	55.8
SAN	**46.6**	**56.5**

Table 5: Test performance on the adversarial SQuAD dataset in F1 score.

Conclusion

- Neural approaches to QA = encoding + **reasoning** + decoding
- Learning to reason for KB QA
  - Symbolic: comprehensible but not robust
  - Neural: robust but not comprehensible
  - Hybrid: robust and comprehensible
- Learning to reason for Text QA / MRC
  - Need better tasks / datasets ! – MS MARCO?
  - ReasoNet: Learning when to step via RL
  - SAN: stochastic prediction dropout