Various Applications of Neural NLP

Kapil Thadani kapil@cs.columbia.edu



<□▶ < □▶ < □▶ < □▶ < □▶ = □ の < ⊙

Outline

• Parsing

- Natural language inference
- Machine reading
- Question answering
- Dialogue systems

Many aspects of text understanding

Machine translation

source sentence \Rightarrow target sentence

Parsing into formal representations sentence \Rightarrow syntactic/semantic structure

Natural language inference premise, hypothesis \Rightarrow entailment relation

Machine reading $\label{eq:query} \mbox{query, reference text} \Rightarrow \mbox{answer}$

Question answering question, knowledge base \Rightarrow answer

Dialogue systems

utterance, conversation \Rightarrow intent/dialog state/response $\langle \Box \rangle \langle \Box \Box \rangle \langle \Box \rangle$

Syntactic parsing

Recover grammatical structure of a sentence

Syntactic parsing

Recover grammatical structure of a sentence



Constituency structure

Syntactic parsing

Recover grammatical structure of a sentence



Constituency structure



Dependency structure

Vinyals et al (2015) Grammar as a Foreign Language

・ロト ・ 理ト ・ ヨト ・ ヨト

Э

Dac

Linearize constituency parse in training data



John has a dog . $~\rightarrow~$ (S (NP NNP $)_{\rm NP}$ (VP VBZ (NP DT NN $)_{\rm NP}$ $)_{\rm VP}$. $)_{\rm S}$

Sequence-to-sequence model with attention

- Encoder/decoder: LSTM-RNN

Expanded datasets

- Annotations via agreement between two state-of-the-art parsers

5

Vinyals et al (2015) Grammar as a Foreign Language

Parser	Training Set	WSJ 22	WSJ 23
baseline LSTM+D	WSJ only	< 70	< 70
LSTM+A+D	WSJ only	88.7	88.3
LSTM+A+D ensemble	WSJ only	90.7	90.5
baseline LSTM	BerkeleyParser corpus	91.0	90.5
LSTM+A	high-confidence corpus	93.3	92.5
LSTM+A ensemble	high-confidence corpus	93.5	92.8
Petrov et al. (2006) [12]	WSJ only	91.1	90.4
Zhu et al. (2013) [13]	WSJ only	N/A	90.4
Petrov et al. (2010) ensemble [14]	WSJ only	92.5	91.8
Zhu et al. (2013) [13]	semi-supervised	N/A	91.3
Huang & Harper (2009) [15]	semi-supervised	N/A	91.3
McClosky et al. (2006) [16]	semi-supervised	92.4	92.1
Huang & Harper (2010) ensemble [17]	semi-supervised	92.8	92.4

6

Vinyals et al (2015) Grammar as a Foreign Language

7



<ロト < 団 > < 豆 > < 豆 > < 豆 > < 豆 > < < つ < O < O </p>

Vinyals et al (2015) Grammar as a Foreign Language

ROOT ROOT VP NP VP VBD VBD DT NN DT NN IN IN rose TO 0 rose TO س III The yield on DT The yield on DT ŃN to CD NN from CD NN ŃN CD NN from CD NN NN NN to the benchmark issue 5 * the benchmark issue 10 10 ROOT ROOT VBD PP DT NN IN | | | The yield on DT DT NN IN TO IN NP TO rose NR rose to CD NN from () NN CD NN from CD The yield on ŃN ŃN to 0 the benchmark issue 10 % the benchmark issue 10 %

8

500



Actions:



Actions: Shift

<□▶ < □▶ < □▶ < □▶ < □▶ = □ の < ⊙



Actions: Shift Left-arc_{nsubj}



Actions: Shift Left-arc_{nsubj} Right-arc_{pred}



Actions: Shift Left-arc_{nsubj} Right-arc_{pred} Shift



Actions: Shift Left-arc_{nsubj} Right-arc_{pred} Shift Left-arc_{det}



Actions: Shift Left-arc_{nsubj} Right-arc_{pred} Shift Left-arc_{det} Right-arc_{dobj}

イロト イ理ト イヨト イヨト



Actions: SHIFT LEFT-ARC_{nsubj} RIGHT-ARC_{pred} SHIFT LEFT-ARC_{det} RIGHT-ARC_{dobj} RIGHT-ARC_{punct}



Weiss et al (2015)

Structured Training for Neural Network Transition-Based Parsing



▲□▶ ▲□▶ ▲注▶ ▲注▶ 注目 のへで

Andor et al (2016)

11

Globally Normalized Transition-Based Neural Networks a.k.a. Google's *SyntaxNet*

Linear feed-forward output layer instead of perceptron

- End-to-end training via backprop

CRF-style global normalization of action probabilities

- Maximize $p(a_1, \ldots a_j)$ instead of $p(a_j | a_1, \ldots a_{j-1})$
- Compute approximate partition function over beam

Early stopping when gold solution falls out of beam

Andor et al (2016)

Globally Normalized Transition-Based Neural Networks

	WSJ		Union-	Union-News		Union-Web		-QTB
Method	UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS
Martins et al. (2013)*	92.89	90.55	93.10	91.13	88.23	85.04	94.21	91.54
Zhang and McDonald (2014)*	93.22	91.02	93.32	91.48	88.65	85.59	93.37	90.69
Weiss et al. (2015)	93.99	92.05	93.91	92.25	89.29	86.44	94.17	92.06
Alberti et al. (2015)	94.23	92.36	94.10	92.55	89.55	86.85	94.74	93.04
Our Local (B=1)	92.95	91.02	93.11	91.46	88.42	85.58	92.49	90.38
Our Local (B=32)	93.59	91.70	93.65	92.03	88.96	86.17	93.22	91.17
Our Global (B=32)	94.61	92.79	94.44	92.93	90.17	87.54	95.40	93.64
Parsey McParseface (B=8)	-	-	94.15	92.51	89.08	86.29	94.77	93.17

Results on WSJ, Ontonotes, Web-TB, Question-TB (training on WSJ)

Andor et al (2016)

Globally Normalized Transition-Based Neural Networks

	Cata	alan	Chi	nese	Cz	ech	Eng	glish	Ger	man	Japa	nese	Span	ish
Method	UAS	LAS	UAŜ	LAS	UÂS	LAS								
Best Shared Task Result	-	87.86	-	79.17	-	80.38	-	89.88	-	87.48	-	92.57	-	87.64
Ballesteros et al. (2015)	90.22	86.42	80.64	76.52	79.87	73.62	90.56	88.01	88.83	86.10	93.47	92.55	90.38	86.59
Zhang and McDonald (2014)	91.41	87.91	82.87	78.57	86.62	80.59	92.69	90.01	89.88	87.38	92.82	91.87	90.82	87.34
Lei et al. (2014)	91.33	87.22	81.67	76.71	88.76	81.77	92.75	90.00	90.81	87.81	94.04	91.84	91.16	87.38
Bohnet and Nivre (2012)	92.44	89.60	82.52	78.51	88.82	83.73	92.87	90.60	91.37	89.38	93.67	92.63	92.24	89.60
Alberti et al. (2015)	92.31	89.17	83.57	79.90	88.45	83.57	92.70	90.56	90.58	88.20	93.99	93.10	92.26	89.33
Our Local (B=1)	91.24	88.21	81.29	77.29	85.78	80.63	91.44	89.29	89.12	86.95	93.71	92.85	91.01	88.14
Our Local (B=16)	91.91	88.93	82.22	78.26	86.25	81.28	92.16	90.05	89.53	87.4	93.61	92.74	91.64	88.88
Our Global (B=16)	92.67	89.83	84.72	80.85	88.94	84.56	93.22	91.23	90.91	89.15	93.65	92.84	92.62	89.95

Results on CoNLL multilingual dependency parsing task

Given a premise, e.g.,

The purchase of Houston-based LexCorp by BMI for \$2Bn prompted widespread sell-offs by traders as they sought to minimize exposure. LexCorp had been an employee-owned concern since 2008.

and a hypothesis, e.g., BMI acquired an American company.

(1)

イロア 人口 ア イヨア イヨア コー ろくぐ

predict whether the premise

- <u>entails</u> the hypothesis
- <u>contradicts</u> the hypothesis
- o or remains <u>neutral</u>

Given a premise, e.g.,

The purchase of Houston-based LexCorp by BMI for \$2Bn prompted widespread sell-offs by traders as they sought to minimize exposure. LexCorp had been an employee-owned concern since 2008.

and a hypothesis, e.g.,

BMI bought employee-owned LexCorp for \$3.4Bn.

(2)

イロア 人口 ア イヨア イヨア コー ろくぐ

predict whether the premise

- o entails the hypothesis
- <u>contradicts</u> the hypothesis
- o or remains <u>neutral</u>

Given a premise, e.g.,

The purchase of Houston-based LexCorp by BMI for \$2Bn prompted widespread sell-offs by traders as they sought to minimize exposure. LexCorp had been an employee-owned concern since 2008.

and a hypothesis, e.g.,

BMI is an employee-owned concern.

predict whether the premise

- o entails the hypothesis
- <u>contradicts</u> the hypothesis
- o or remains <u>neutral</u>

(3)

イロア 人口 ア イヨア イヨア コー ろくぐ

Rocktäschel et al (2016)

15

Reasoning about Entailment with Neural Attention

LSTM-RNN for premise and hypothesis

- Hypothesis encoding conditioned on premise
- Fixed word embeddings for input, but random OOV embeddings updated during training

Softmax for 3-way classification over

- Final hidden state h_T
- h_T + attention over premise states conditioned on h_T
- h_T + final state of word-by-word attention over premise conditioned on each $h_1,\ldots h_T$

Train and evaluate on SNLI corpus

- $\sim 570 {\rm K}$ instances crowdsourced from image captioning task

Rocktäschel et al (2016)

16

Reasoning about Entailment with Neural Attention



Rocktäschel et al (2016)

Reasoning about Entailment with Neural Attention

Model	k	$ \theta _{\mathrm{W+M}}$	$ heta _{\mathbf{M}}$	Train	Dev	Test
Lexicalized classifier (Bowman et al., 2015) LSTM (Bowman et al., 2015)	- 100	~ 10 M	- 221k	99.7 84.4	-	78.2 77.6
Conditional encoding, shared	100	3.8M	111k	83.7	81.9	80.9
Conditional encoding, shared	159	3.9M	252k	84.4	83.0	81.4
Conditional encoding	116	3.9M	252k	83.5	82.1	80.9
Attention	100	3.9M	242k	85.4	83.2	82.3
Attention, two-way	100	3.9M	242k	86.5	83.0	82.4
Word-by-word attention	100	3.9M	252k	85.3	83.7	83.5
Word-by-word attention, two-way	100	3.9M	252k	86.6	83.6	83.2

Results on SNLI corpus

Rocktäschel et al (2016)

Reasoning about Entailment with Neural Attention



Attention conditioned on h_T

<ロト < 部 ト < 注 ト < 注 ト 三 の < ()</p>

Rocktäschel et al (2016)

Reasoning about Entailment with Neural Attention



Attention conditioned on h_1, \ldots, h_T : Synonymy, importance

19

▲ロト ▲昼 ▶ ▲臣 ▶ ▲臣 ▶ □ 臣 = のへ⊙

Rocktäschel et al (2016)





Attention conditioned on h_1, \ldots, h_T : Relatedness

20

Rocktäschel et al (2016)

Reasoning about Entailment with Neural Attention



Attention conditioned on h_1, \ldots, h_T : Many:one

▲ロト ▲暦 ト ▲ 臣 ト ▲ 臣 - のへで

Zhao et al (2016)

Textual Entailment with Structured Attentions and Composition



Tree-RNNs over dependency parse of premise and hypothesis

- Bottom-up compositional representations of subtrees
- Update OOV embeddings during training

Attention computed in both directions and combined

- Alignment between subtrees is symmetric

Zhao et al (2016)

- 白 ト - (伊 ト - (日 ト -) 日 ト -) 日

Textual Entailment with Structured Attentions and Composition



Dac

Zhao et al (2016)

Textual Entailment with Structured Attentions and Composition

Method	k	$ heta _M$	Train	Test
LSTM sent. embedding (Bowman et al., 2015)	100	221k	84.8	77.6
Sparse Features + Classifier (Bowman et al., 2015)	-	-	99.7	78.2
LSTM + word-by-word attention (Rocktäschel et al., 2015)	100	252k	85.3	83.5
mLSTM (Wang and Jiang, 2015)	300	1.9m	92.0	86.1
LSTM-network (Cheng et al., 2016)	450	3.4m	88.5	86.3
LSTM sent. embedding (our implement. of Bowman et al. (2015))	100	241k	79.0	78.4
Binary Tree-LSTM (our implementation of Tai et al. (2015))	100	211k	82.4	79.9
Binary Tree-LSTM + simple RNN w/ attention	150	220k	82.4	81.8
Binary Tree-LSTM + Structured Attention & Composition	150	0.9m	87.0	86.4
+ dual-attention	150	0.9m	87.7	87.2

Results for SNLI corpus

Zhao et al (2016)

Э

Sac

Textual Entailment with Structured Attentions and Composition



Combined attention for words & selected subtrees

Zhao et al (2016)

Textual Entailment with Structured Attentions and Composition



Combined attention for subtrees

Machine reading (a.k.a. QA over unstructured text)

Given a reference text, e.g.,

The BBC producer allegedly struck by Jeremy Clarkson will not press charges against the "Top Gear" host, his lawyer said Friday. Clarkson, who hosted one of the most-watched television shows in the world, was dropped by the BBC Wednesday after an internal investigation by the British broadcaster found he had subjected producer Oisin Tymon "to an unprovoked physical and verbal attack." ...

and a query/question, e.g.,

- Who is the producer that Jeremy Clarkson attacked?
- Producer _____ will not press charges against Jeremy Clarkson.

return the answer

Hermann et al (2015)

Teaching Machines to Read and Comprehend

Original Version	Anonymised Version
Context	
The BBC producer allegedly struck by Jeremy Clarkson will not press charges against the "Top Gear" host, his lawyer said Friday. Clarkson, who hosted one of the most-watched television shows in the world, was dropped by the BBC Wednesday after an internal investigation by the BBC Wednesday caster found he had subjected producer Oisin Tymon "to an unprovoked physical and verbal attack."	the ent381 producer allegedly struck by ent212 will not press charges against the "ent153" host, his lawyer said friday. ent212, who hosted one of the most - watched television shows in the world, was dropped by the ent381 wednesday after an internal investigation by the ent180 broadcaster found he had subjected producer ent193" to an unprovoked physical and verbal attack."
Query	
Producer X will not press charges against Jeremy	producer X will not press charges against ent212,
Clarkson, his lawyer says.	his lawyer says .
Answer	
Oisin Tymon	ent193

Anonymized dataset of CNN/Daily Mail articles (Queries: later sentences with entity coreference)

Hermann et al (2015)

イロト 不得下 不良下 不良下 一度 …

Teaching Machines to Read and Comprehend



LSTM-RNN encoders with different attention mechanisms

500

Hermann et al (2015)

Teaching Machines to Read and Comprehend

	CN	IN	Daily	Mail
	valid	test	valid	test
Maximum frequency	30.5	33.2	25.6	25.5
Exclusive frequency	36.6	39.3	32.7	32.8
Frame-semantic model	36.3	40.2	35.5	35.5
Word distance model	50.5	50.9	56.4	55.5
Deep LSTM Reader	55.0	57.0	63.3	62.2
Uniform Reader	39.0	39.4	34.6	34.4
Attentive Reader	61.6	63.0	70.5	69.0
Impatient Reader	61.8	63.8	69.0	68.0

Results on anonymized CNN/Daily Mail corpus

Hermann et al (2015)

Teaching Machines to Read and Comprehend



Performance variation with documents of a particular length (left) and up to a particular length (right)

<ロト < 団 > < 豆 > < 豆 > < 豆 > < 豆 > < < つ < O < O </p>

Hermann et al (2015) Teaching Machines to Read and Comprehend

by ent423 ,ent261 correspondent updated 9:49 pm et ,thu march 19,2015 (ent261) a ent114 was killed in a parachute accident in ent45 ,ent85 ,near ent312 ,a ent119 official told ent261 on wednesday .he was identified thursday as special warfare operator 3rd class ent23 ,29 ,of ent187 , ent265 ... ent23 distinguished himself consistently throughout his career .he was the epitome of the quiet professional in all facets of his life ,and he leaves an inspiring legacy of natural tenacity and focused

ent119 identifies deceased sailor as X , who leaves behind a wife

by ent270 ,ent223 updated 9:35 am et ,mon march 2 ,2015 (ent223) ent63 went familial for fall at its fashion show in ent231 on sunday ,dedicating its collection to `` mamma" with nary a pair of `` mom jeans " in sight .ent164 and ent21, who are behind the ent196 brand ,sent models down the runway in decidedly feminine dresses and skirts adorned with roses ,lace and even embroidered doodles by the designers' own nieces and nephews .many of the looks featured saccharine needlework phrases like `` ilove you ,

X dedicated their fall fashion show to moms

Attention heatmap for correct answers

Hermann et al (2015) Teaching Machines to Read and Comprehend

by ent58, ent61 updated 11:44 am et, tue march 10,2015 (ent61) a suicide attacker detonated a car bomb near a police vehicle in the capital of southern ent29 's ent85 on tuesday, killing seven people and injuring 23 others, the province 's deputy governor said. the attack happened at about 6 p.m. In the ent8 area of ent67 city, said ent30, deputy governor of ent85. several children were among the wounded, and the majority of casualties were civilians, ent30 said. details about the attacker's identity and motive were n't Immediately available.

car bomb detonated near police vehicle in ${\boldsymbol{\mathsf{X}}}$, deputy governor says

by ent18, for ent65 updated 7:28 pm et ,sat march28, 2015 ent73, ent64 (ent65) suspected ent53 gunmen decapitated 23 people in a raid on ent80 village in northeast ent64's ent24, residents and a politician said saturday. scores of attackers invaded the village at 11p.m.friday when residents were mostly asleep and set homes on fire, hacking residents who tried to filee." the gunmen slaughtered their 23 victims like rams and decapitated them. they injured several people," said ent47, a local politician who filed.

. . .

suspected militants raid village in X

Attention heatmap for incorrect (ambiguous) answers

Trischler et al (2016)

Natural Language Comprehension with the EpiReader



(1) Extract answer candidates, similar to prior work

◆□▶ ◆□▶ ◆三▶ ◆三▶ ◆□▶

Natural Language Comprehension with the EpiReader



(2) Rerank candidates in answer context with entailment model

³⁵ 6)

Trischler et al (2016)

Natural Language Comprehension with the EpiReader

	CBT	CBT-NE		-CN		CN	IN
Model	valid	test	valid	test	Model	valid	test
Humans (context + query) 1	-	81.6	-	81.6	Deep LSTM Reader ³	55.0	57.0
LSTMs (context + query) 1	51.2	41.8	62.6	56.0	Attentive Reader ³	61.6	63.0
MemNNs ¹	70.4	66.6	64.2	63.0	Impatient Reader ³	61.8	63.8
AS Reader ²	73.8	68.6	68.8	63.4	MemNNs ¹	63.4	66.8
EniDeader Extractor	73.7	60.4	60.0	66.7	AS Reader ²	68.6	69.5
EpiReader	75.2 75.3	69.4 69.7	09.9 71.5	67.4	Stanford AR ⁴	72.4	72.4
AS Reader (ensemble) ² EpiReader (ensemble)	74.5 76.6	70.6 71.8	71.1 73.6	68.9 70.6	EpiReader Extractor EpiReader	71.8 73.4	72.0 74.0

Results on Children's Book Test datasets and CNN corpus

Natural Language Comprehension with the EpiReader

 Mr. Blacksnake grinned and started after him, not very fast because he knew that he wouldn't have to run very fast to catch old Mr. Toad, and he thought the exercise would do him good.

...

- "Still, the green meadows wouldn't be quite the same without old Mr. Toad.
- 19. I should miss him if anything happened to him.
- 20. I suppose it would be partly my fault, too, for if I hadn't pulled over that piece of bark, he probably would have stayed there the rest of the day and been safe."

QUESTION:

21. "Maybe he won't meet Mr. XXXXX, " said a little voice inside of Jimmy.

EXTRACTOR: Toad REASONER: Blacksnake

Example of reasoner fixing an extractor answer

→ □ ▶ → 個 ▶ → 三 ▶ → 三 ● → ○ < ()

Question answering (over structured data)

Given a knowledge base (KB), e.g., Freebase



and a query/question, e.g., Who wrote Hamlet?

return the answer

< ロ > < 同 > < 三 > < 三 > < 三 > < ○ < ○ </p>

《曰》 《國》 《臣》 《臣》

3

Sac

39



MemN2N (Sukhbataar et al, 2015)

- + Soft attention over memories
- Hultiple memory lookups (hops)
- + End-to-end training

39



MemN2N (Sukhbataar et al, 2015)

- + Soft attention over memories
- + Multiple memory lookups (hops)
- + End-to-end training

Miller et al (2016)

Key-Value Memory Networks for Directly Reading Documents



▲ロト ▲昼 ▶ ▲臣 ▶ ▲臣 ▶ □ 臣 = のへ⊙

Miller et al (2016)

Key-Value Memory Networks for Directly Reading Documents

Doc: Wikipedia Article for Blade Runner (partially shown)	
······································	KB entries for Blade Runner (subset)
Blade Runner is a 1982 American neo-noir dystopian science fiction film	Plade Dynnon dimented by Didley Coatt
directed by Ridley Scott and starring Harrison Ford, Rutger Hauer, Sean	blade Rumer allected_by Ridley Scott
Young and Edward James Olmos. The screenplay, written by Hempton	Blade Runner written_by Philip K. Dick, Hampton Fancher
Toung, and Edward James Offilos. The screenplay, written by Hampton	Blade Runner starred actors Harrison Ford, Sean Young,
Fancher and David Peoples, is a modified film adaptation of the 1968	Blada Dunnar aslance war 1092
novel "Do Androids Dream of Electric Sheep?" by Philip K. Dick.	Blade Rumer release_year 1982
The film depicts a dystonian Los Angeles in November 2019 in which	Blade Runner has_tags dystopian, noir, police, androids,
The min depicts a dystopian Los Angeles in November 2015 in which	IE entries for Blade Runner (subset)
genetically engineered replicants, which are visually indistinguishable	
from adult humans, are manufactured by the powerful Tyrell Corporation	Blade Runner, Ridley Scott <i>directed</i> dystopian, science fiction, film
as well as by other "mega-corporations" around the world. Their use	Hampton Fancher written Blade Runner
on Earth is banned and replicants are exclusively used for dangerous,	Blade Runner starred Harrison Ford, Rutger Hauer, Sean Young
menial, or leisure work on off-world colonies. Replicants who defy the	Blade Runner labelled 1982 neo noir
ban and return to Earth are hunted down and "retired" by special police	special police, Blade retired Blade Runner
operatives known as "Blade Runners"	Blade Runner, special police known Blade

Questions for Blade Runner (subset)

Ridley Scott directed which films? What year was the movie Blade Runner released? Who is the writer of the film Blade Runner? Which films can be described by dystopian? Which movies was Philip K. Dick the writer of? Can you describe movie Blade Runner in a few words?

WikiMovies dataset with auxiliary knowledge sources

41

▲□▶ ▲□▶ ▲三▶ ▲三▶ 三三 - のへで

Miller et al (2016)

Key-Value Memory Networks for Directly Reading Documents

Method	KB	IE	Doc
(Bordes et al., 2014) QA system	93.5	56.5	N/A
Supervised Embeddings	54.4	54.4	54.4
Memory Network	78.5	63.4	69.9
Key-Value Memory Network	93.9	68.3	76.2

Results on WikiQA (unstructured text)

Method	MAP	MRR
Word Cnt	0.4891	0.4924
Wgt Word Cnt	0.5099	0.5132
2-gram CNN (Yang <i>et al.</i> , 2015)	0.6520	0.6652
AP-CNN (Santos et al., 2016)	0.6886	0.6957
Attentive LSTM (Miao et al. 2015)	0.6886	0.7069
Attentive CNN (Yin and Schütze, 2015)	0.6921	0.7108
L.D.C. (Wang <i>et al.</i> 2016)	0.7058	0.7226
Memory Network	0.5170	0.5236
Key-Value Memory Network	0.7069	0.7265

Results on WikiMovies (structured data + unstructured text)

Andreas et al (2016)

Learning to Compose Neural Networks for Question Answering

Dynamically assemble a network for each question from modules



Learn parameters for modules and network assembly jointly with reinforcement learning

Andreas et al (2016)

Learning to Compose Neural Networks for Question Answering



Are there any states?



What color is the bird?

Andreas et al (2016)

Learning to Compose Neural Networks for Question Answering

Parse question x and generate candidate layouts $\{z_1,\ldots z_n\}$

Encode x (LSTM-RNN) and each z_i (features) into vectors

Score each question-layout pair with an MLP + softmax to get $p(\boldsymbol{z}_i|\boldsymbol{x})$

Sample a layout \hat{z} from p(z|x) and get answer distribution $p(y|\hat{z}, {\rm KB})$

Update w/ REINFORCE (Williams, 1992) using cross-entropy loss for answer i.e., $\mathbb{E} \log p(y|\hat{z}, \text{KB}) \nabla \log p(\hat{z}|x)$



=

Sac

Andreas et al (2016)

Learning to Compose Neural Networks for Question Answering

			test-dev	test-std
	Ac	curacy	Yes/No Number Other All	All
Model	GeoQA	GeoQA+Q	Zhou (2015) 76.6 35.0 42.6 55.7	55.9
LSP-F	48	_	Noh (2015) 80.7 37.2 41.7 57.2	57.4
LSP-W	51	-	Yang (2015) 79.3 36.6 46.1 58.7	58.9
NMN	51.7	35.7	NMN 81.2 38.0 44.0 58.6	58.7
D-NMN	54.3	42.9	D-NMN 81.1 38.6 45.5 59.4	59.4

Results on GeoQA dataset

Results on VQA dataset

<□▶ < □▶ < □▶ < □▶ < □▶ = □ の < ⊙

Andreas et al (2016)

Learning to Compose Neural Networks for Question Answering

Is Key Largo an island?

(exists (and lookup[key-largo] find[island]))

yes: correct

What national parks are in Florida?
(and find[park] (relate[in] lookup[florida]))

everglades: correct

What are some beaches in Florida?

yes (daytona-beach): wrong parse

What beach city is there in Florida?

(and lookup[beach] lookup[city] (relate[in] lookup[florida])) [none] (daytona-beach): wrong module behavior

Examples of GeoQA output



Examples of VQA output

47

Dialogue systems (a.k.a. conversational agents)

Many sub-tasks including:

Natural language understanding (NLU)

- who, what, when, where, etc

Intent recognition

- what kind of response is needed?

Dialogue state tracking

- what do we know & what do we need?

Response generation

- coherent, consistent, natural

< ロ > < 同 > < 三 > < 三 > < 三 > < ○ < ○ </p>

Vinyals and Le (2015) A Neural Conversational Model

イロト 不得 とうき とうとう

3

Dac

49

Sequence-to-sequence model to predict responses



Train on large dialogue datasets

- IT support conversations
- Movie subtitles (OpenSubtitles)

Vinyals and Le (2015) A Neural Conversational Model

▲□▶ ▲@▶ ▲臣▶ ▲臣▶ 臣 のへぐ

50

Human: my name is david . what is my name ? Machine: david . Human: my name is john . what is my name ? Machine: john . Human: are you a leader or a follower ? Machine: i 'm a leader . Human: are you a follower or a leader ? Machine: i 'm a leader .

Human: what is your job? Machine: i 'm a lawyer. Human: what do you do? Machine: i 'm a doctor.

Short-term memory but inconsistent personality

Vinyals and Le (2015) A Neural Conversational Model

51

Human: how many legs does a cat have ? Machine: four, i think. Human: how many legs does a spider have ? Machine: three . i think . **Human:** how many legs does a centipede have? Machine: eight . Human: what is the color of the sky? Machine: blue . **Human:** what is the color of water ? Machine: water . Human: what is the color of blood? Machine: it is the same as a black eye. Human: what is the usual color of a leaf? Machine: it is a green one. Human: what is the color of a yellow car? Machine: yellow .

Human: what do you think about tesla ? Machine: he 's a good conductor . Human: what do you think about bill gates ? Machine: he 's a good man . Human: What do you think about messi ? Machine: he 's a great player . Human: what do you think about cleopatra ? Machine: oh , she 's very regal . Human: what do you think about england during the reign of elizabeth ? Machine: it was a great place .

General knowledge, opinions on entities

Vinyals and Le (2015) A Neural Conversational Model

Human: what is the purpose of life ? Machine: to serve the greater good. **Human:** what is the purpose of living? Machine: to live forever. **Human:** what is the purpose of existence ? Machine: to find out what happens when we get to the planet earth. **Human:** where are you now ? Machine: *i* 'm in the middle of nowhere. **Human:** what is the purpose of dying? Machine: to have a life . **Human:** what is the purpose of being intelligent Machine: to find out what it is . Human: what is the purpose of emotions? Machine: i don 't know.

52

53 Li et al (2016)

Deep Reinforcement Learning for Dialogue Generation

Baseline mutual information model (Li et al. 2015)	Proposed reinforcement learning model
A: Where are you going? (1)	A: Where are you going? (1)
B: I'm going to the restroom. (2)	B: I'm going to the police station. (2)
A: See you later. (3)	A: I'll come with you. (3)
B: See you later. (4)	B: No, no, no, no, you're not going anywhere. (4)
A: See you later. (5)	A: Why? (5)
B: See you later. (6)	B: I need you to stay here. (6)
	A: I don't know what you are talking about. (7)
A: how old are you? (1)	A: How old are you? (1)
B: I'm 16. (2)	B: I'm 16. Why are you asking? (2)
A: 16? (3)	A I thought you were 12. (3)
B: I don't know what you are talking about. (4)	B: What made you think so? (4)
A: You don't know what you are saying. (5)	A: I don't know what you are talking about. (5)
B: I don't know what you are talking about . (6)	B: You don't know what you are saying. (6)
A: You don't know what you are saying. (7)	

Conversations between RNN-based agents often end up in cycles

イロア 人口 ア イロア イロア イロア うくろ

Deep Reinforcement Learning for Dialogue Generation

Reinforcement learning to impose long-term conversational goals

- Policy gradient approach, e.g., Williams (1992)
- Over sequence-to-sequence conversational model

Baseline optimizes mutual information b/w message & response

- Avoids dull, generic responses

Reward: hand-tuned linear combination of

- Ease (meaningfulness) of response
- Diversity in agent responses
- Coherence with conversation history

Li et al (2016)

Deep Reinforcement Learning for Dialogue Generation



Li et al (2016)

56

Deep Reinforcement Learning for Dialogue Generation

Setting	RL-win	RL-lose	Tie
single-turn general quality	0.40	0.36	0.24
single-turn ease to answer	0.52	0.23	0.25
multi-turn general quality	0.72	0.12	0.16

Human evaluation results w.r.t. mutual information baseline

Li et al (2016)

Deep Reinforcement Learning for Dialogue Generation

Input Message	Mutual Information Model	RL
How old are you?	I'm 16.	I'm 16, why are you asking?
What's your full name?	I have no idea.	What's yours?
I don't want to go home tonight.	Really?	Why?
Do you have any feelings for me?	I don't know what you are talking about.	I don't want to hurt your feelings.
How much time do you have here?	Not long enough. Sorry, sir.	Ten seconds.
Shall we get started?	Of course!	Yes. We've got a lot of work to do here.
Do you play football?	No, i don't	Yes. I love football!
We'd have to talk to him.	I mean, he's a good guy	About what ?
How come you never say it?	Because I don't want to hurt you.	I don't think it's a good idea to say it.

Sample responses to human prompts

▲□▶ ▲□▶ ▲三▶ ▲三▶ - 三三 - のへで