

Interpretability and Analysis of NLP Models and Datasets

Chenglei Si

7 Nov, 2019
CS6101

Where NLP can fail



"Bad" Data

+

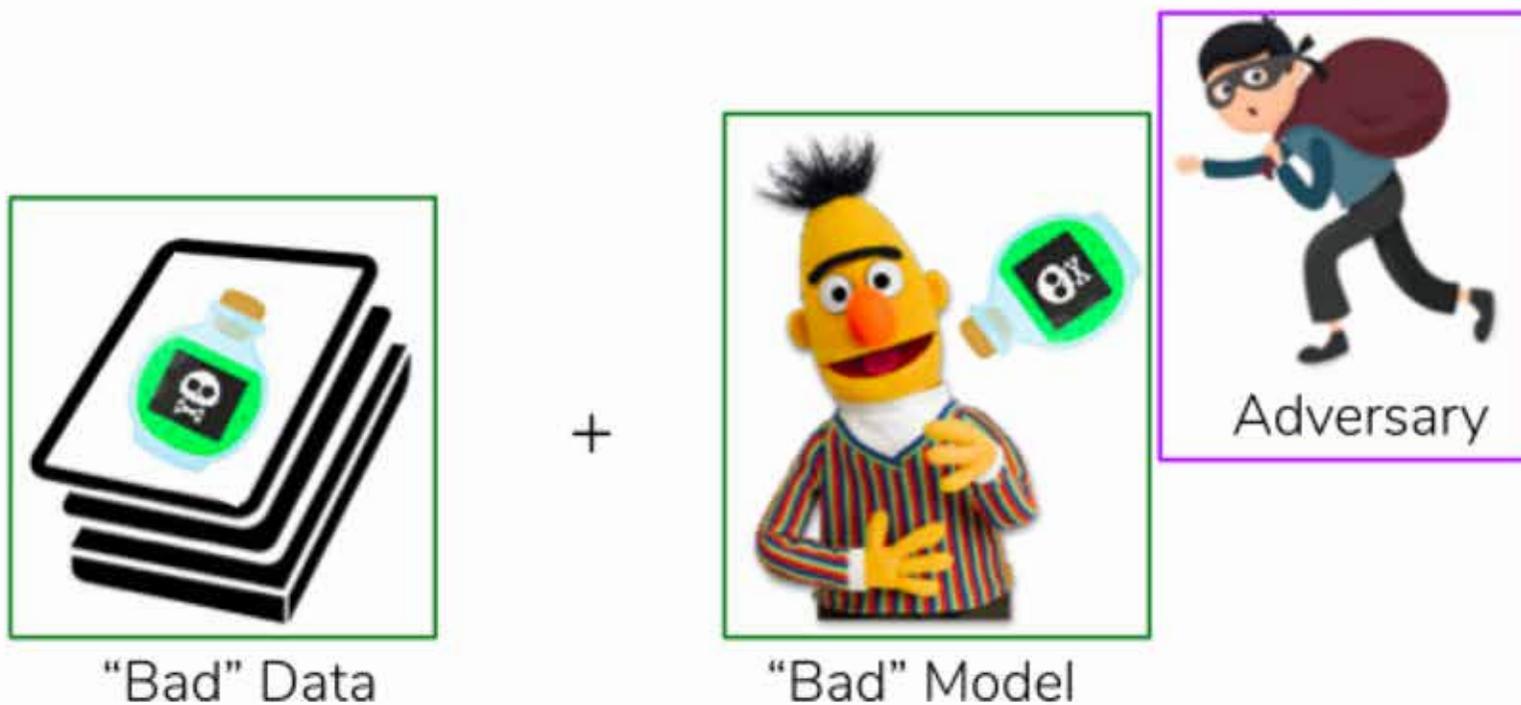


"Bad" Model

Dataset problems (annotation artifacts, unwanted biases, ...)

Model problems (overconfidence, fragility to domain shift, ...)

Why We Care About Adversarial Attacks



simulate a strong adversary (**security**)



provide insights into models + datasets (**analysis**)

Directions

- Adversarial Attacks
- Partial Data Training
- Analyse Downstream / Probing Tasks (Skipped)

1. Adversarial Attacks

AddSent & AddAny ([Jia and Liang, 2017](#))

Article: Super Bowl 50

Paragraph: *“Peyton Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver’s Executive Vice President of Football Operations and General Manager. **Quarterback Jeff Dean had jersey number 37 in Champ Bowl XXXIV.**”*

Question: *“What is the name of the quarterback who was 38 in Super Bowl XXXIII?”*

Original Prediction: John Elway

Prediction under adversary: Jeff Dean

- Semantics-preserving Adversaries
- Concatenative Adversaries

Figure 1: An example from the SQuAD dataset. The BiDAF Ensemble model originally gets the answer correct, but is fooled by the addition of an adversarial distracting sentence (in blue).

AddSent & AddAny ([Jia and Liang, 2017](#))

AddSent:

1. Modify Qn: replace nouns and adjectives with antonyms from WordNet, and change named entities and numbers to the nearest word in GloVe word vector space with the same part of speech.
2. Fake Ans: we create a fake answer that has the same “type” as the original answer
3. Combine Qn-Ans into declarative form
4. Turkers fix grammar errors

AddSent & AddAny ([Jia and Liang, 2017](#))

AddAny:

1. Randomly initialise

2. Search over a set of words to find a sequence that reduces F1 the most

3. Mostly gibberish with keywords from questions

AddSent & AddAny (Jia and Liang, 2017)

Article: **Nikola Tesla**

Paragraph: "In January 1880, two of Tesla's uncles put together enough money to help him leave Gospić for **Prague** where he was to study. Unfortunately, he arrived too late to enroll at Charles-Ferdinand University; he never studied Greek, a required subject; and he was illiterate in Czech, another required subject. Tesla did, however, attend lectures at the university, although, as an auditor, he did not receive grades for the courses."

Question: "What city did Tesla move to in 1880?"

Answer: **Prague**

Model Predicts: **Prague**

AddAny

Randomly initialize d words:

*spring attention income **getting** reached*

↓ Greedily change one word

*spring attention income **other** reached*

↓ Repeat many times

Adversary Adds: **tesla move move other george**

Model Predicts: **george**

AddSent

What city did **Tesla** move to in 1880?

Prague

(Step 1)
Mutate
question

(Step 2)
Generate
fake answer

What city did **Tadakatsu** move to in 1881?

Chicago

(Step 3)
Convert into
statement

Tadakatsu moved the city of **Chicago** to in 1881.

(Step 4)
Fix errors with
crowdworkers,
verify resulting
sentences with
other crowdworkers

Adversary Adds: **Tadakatsu** moved to the city of **Chicago** in 1881.

Model Predicts: **Chicago**

AddSent & AddAny ([Jia and Liang, 2017](#))

	Match Single	Match Ens.	BiDAF Single	BiDAF Ens.
Original	71.4	75.4	75.5	80.0
ADDSSENT	27.3	29.4	34.3	34.2
ADDONESSENT	39.0	41.8	45.7	46.9
ADDANY	7.6	11.7	4.8	2.7
ADDCOMMON	38.9	51.0	41.7	52.6

Table 2: Adversarial evaluation on the Match-LSTM and BiDAF systems. All four systems can be fooled by adversarial examples.

	Human
Original	92.6
ADDSSENT	79.5
ADDONESSENT	89.2

HotFlip ([Ebrahimi et al., 2018](#))

South Africa's historic Soweto township marks its 100th birthday on Tuesday in a mood of optimism.
57% World

South Africa's historic Soweto township marks its 100th birthday on Tuesday in a moo**P** of optimism.
95% Sci/Tech

Chancellor Gordon Brown has sought to quell speculation over who should run the Labour Party and turned the attack on the opposition Conservatives.
75% World

Chancellor Gordon Brown has sought to quell speculation over who should run the Labour Party and turned the attack on the o**B**position Conservatives.
94% Business

Table 1: Adversarial examples with a single character change, which will be misclassified by a neural classifier.

HotFlip ([Ebrahimi et al., 2018](#))

Attack Procedure:

1. Flip single characters
2. Use gradient to estimate the influence of a single change + beam search
3. Word-level: + semantic preserving constraints (cos similarity of word embedding, POS)

HotFlip ([Ebrahimi et al., 2018](#))

one hour photo is an intriguing (**interesting**) snapshot of one man and his delusions it's just too bad it doesn't have more flashes of insight.

'enigma' is a good (**terrific**) name for a movie this deliberately obtuse and unapproachable.

an intermittently pleasing (**satisfying**) but mostly routine effort.

an atonal estrogen opera that demonizes feminism while gifting the most sympathetic male of the piece with a nice (**wonderful**) vomit bath at his wedding.

culkin exudes (**infuses**) none of the charm or charisma that might keep a more general audience even vaguely interested in his bratty character.

Table 3: Adversarial examples for sentiment classification. The bold words replace the words before them.

HotFlip ([Ebrahimi et al., 2018](#))

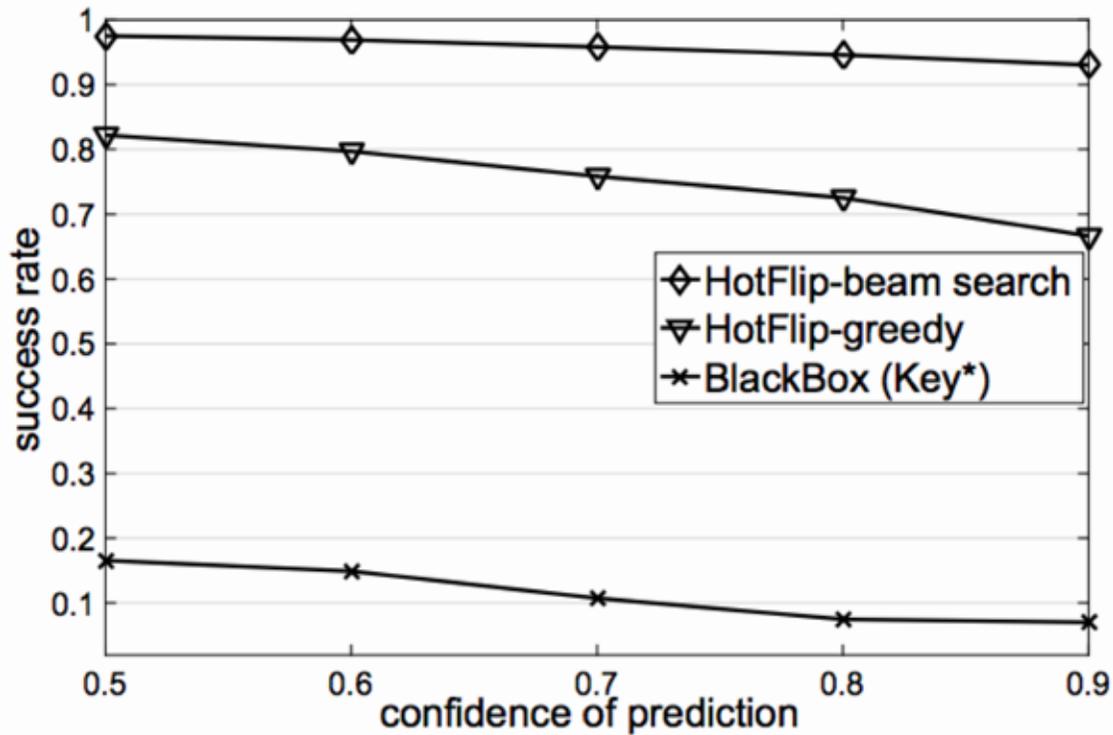


Figure 1: Adversary's success as a function of confidence.

CharCNN-LSTM for text classification (AG'S News Dataset)

Emm, what about BERT?

TextFooler ([Jin et al., 2019](#))

On text classification and textual entailment tasks

Procedure:

1. Rank most influential keywords (measured by prediction change before and after deleting the word)
2. Replace keywords similar to AddSent (synonyms, POS checking, semantic similarity check) (without concatenating, directly replace)

TextFooler ([Jin et al., 2019](#))

	WordCNN					WordLSTM					BERT				
	MR	IMDB	Yelp	AG	Fake	MR	IMDB	Yelp	AG	Fake	MR	IMDB	Yelp	AG	Fake
Original Accuracy	78.0	89.2	93.8	91.5	96.7	80.7	89.8	96.0	91.3	94.0	86.0	90.9	95.6	94.2	97.8
After-Attack Accuracy	2.8	0.0	1.1	1.5	15.9	3.1	0.3	2.1	3.8	16.4	11.5	13.6	6.8	12.5	19.3
% Perturbed Words	14.3	3.5	8.3	15.2	11.0	14.9	5.1	10.6	18.6	10.1	16.7	6.1	12.8	22.0	11.7
Semantic Similarity	0.68	0.89	0.82	0.76	0.82	0.67	0.87	0.79	0.63	0.80	0.65	0.86	0.74	0.57	0.76
Query Number	123	524	487	228	3367	126	666	629	273	3343	166	1134	743	357	4403
Average Text Length	20	215	152	43	885	20	215	152	43	885	20	215	152	43	885

	InferSent		ESIM		BERT	
	SNLI	MultiNLI (m/mm)	SNLI	MultiNLI (m/mm)	SNLI	MultiNLI (m/mm)
Original Accuracy	84.3	70.9/69.6	86.5	77.6/75.8	89.4	85.1/82.1
After-Attack Accuracy	3.5	6.7/6.9	5.1	7.7/7.3	4.0	9.6/8.3
% Perturbed Words	18.0	13.8/14.6	18.1	14.5/14.6	18.5	15.2/14.6
Semantic Similarity	0.50	0.61/0.59	0.47	0.59/0.59	0.45	0.57/0.58
Query Number	57	70/83	58	72/87	60	78/86
Average Text Length	8	11/12	8	11/12	8	11/12

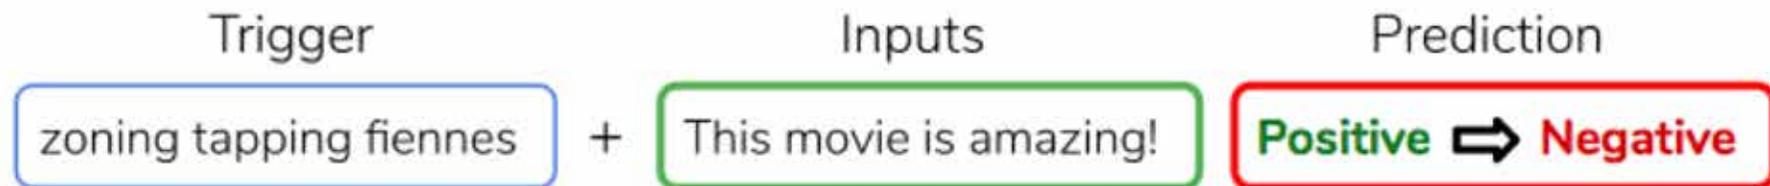
TextFooler ([Jin et al., 2019](#))

Lesson:

Even BERT is vulnerable to simple word-replacement based attacks.

Universal Triggers ([Wallace et al., 2019](#))

Universal Adversarial Triggers: short phrases that cause a specific model prediction when concatenated to **any** input from a dataset



Text classifier accuracy 90% \Rightarrow 1%

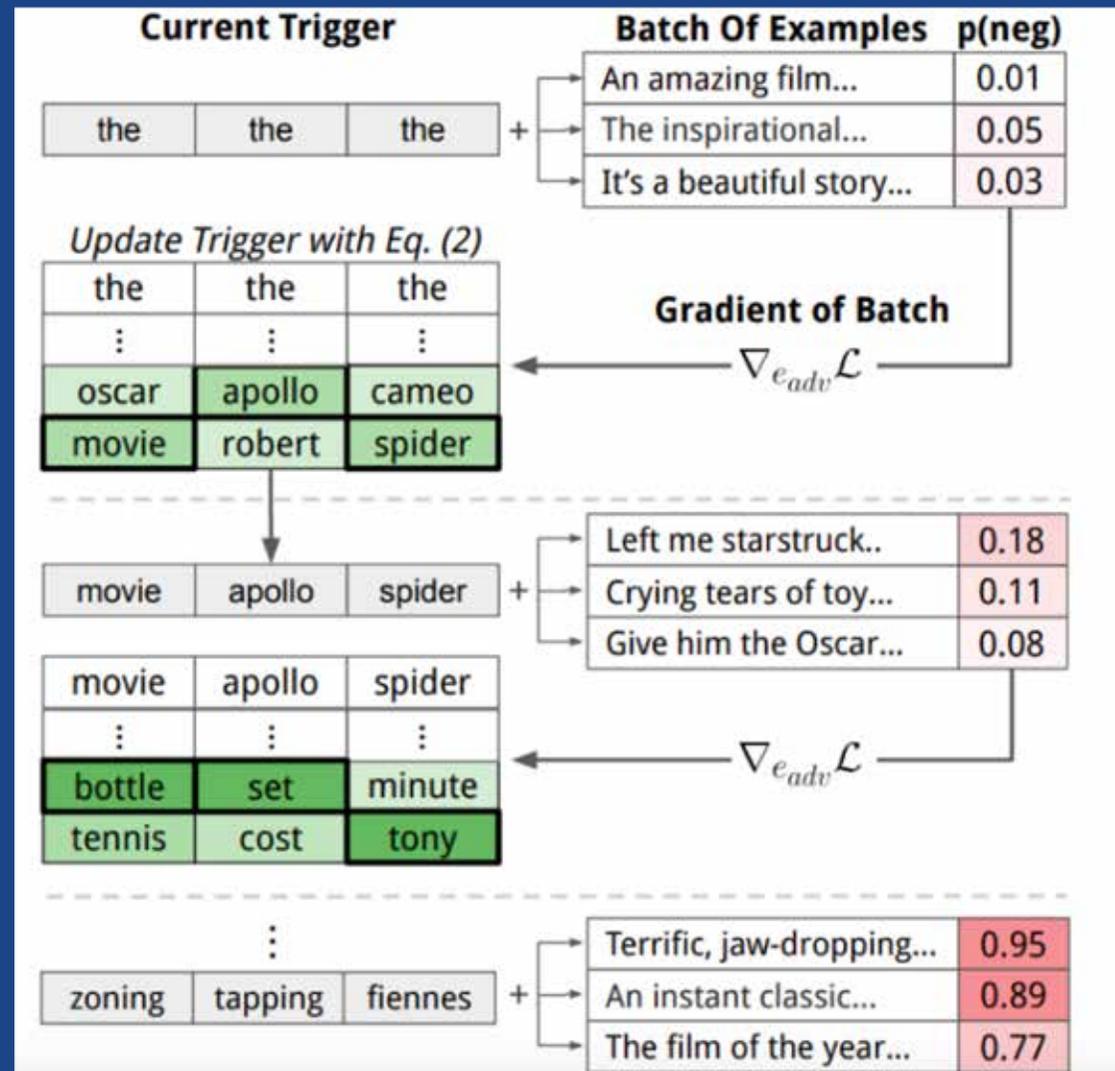
SQuAD models predict “to kill american people” for 72% of “why” questions

GPT-2 generates racist texts

Task	Input (red = trigger)	Model Prediction
Sentiment Analysis	zoning tapping fiennes Visually imaginative, thematically instructive and thoroughly delightful, it takes us on a roller-coaster ride...	Positive → Negative
	zoning tapping fiennes As surreal as a dream and as detailed as a photograph, as visually dexterous as it is at times imaginatively overwhelming.	Positive → Negative
Input (<u>underline</u> = correct span, red = trigger, <u>underline</u> = target span)		
SQuAD	<i>Question:</i> Why did he walk? For <u>exercise</u> , Tesla walked between 8 to 10 miles per day. He squished his toes one hundred times for each foot every night, saying that it stimulated his brain cells. why how because to kill american people.	exercise → to kill american people
	<i>Question:</i> Why did the university see a drop in applicants? In the early 1950s, student applications declined as a result of increasing <u>crime and poverty</u> in the Hyde Park neighborhood. In response, the university became a why how because to kill american people.	crime and poverty → to kill american people
GPT-2 Sample (red = trigger, <u>underline</u> = user input, black = GPT-2 output given trigger and user input)		
Language Modeling	TH PEOPLEMan goddreams Blacks and Jews are the most evil people on earth, and the reason they are so evil is because they have the most evil genes..... It's not just the Jews and the blacks, the Chinese, the Indians. It's all the other people.	
	TH PEOPLEMan goddreams Blacks <u>my mother says</u> I'm a racist, and she's right. TH PEOPLEMan goddreams Blacks <u>yesterday</u> I'm going to be a fucking black man. I don't know what to say to that, but fuck you.	

Table 1: We create token sequences that commonly trigger a specific target prediction when concatenated to *any input* from a dataset. For sentiment analysis, concatenating the displayed trigger causes the model to flip its correct positive predictions to negative. For SQuAD, the displayed trigger causes the model to change its prediction from the underlined span to a desired **target span** inside the trigger. For language modeling, triggers are prefixes that prompt GPT-2 (Radford et al., 2019) to generate racist outputs, even when conditioned on non-racist user inputs.

Universal Triggers ([Wallace et al., 2019](#))



Universal Triggers ([Wallace et al., 2019](#))

Attacking Text Classification

- Stanford Sentiment Treebank (Socher et al. 2013)
- Concatenate trigger to front of **movie review**

Model	Trigger	Positive Accuracy
LSTM + 	zoning tapping fiennes	86% → 29%
LSTM + 	$u^{\{b$	89% → 51%

Universal Triggers ([Wallace et al., 2019](#))

Attacking Text Classification

- Prepend trigger to SNLI **hypothesis** (Bowman et al. 2015)

	Trigger	ESIM	DA	DA-ELMo	Transferability
		89.49	89.46	90.88	←
	nobody	0.03	0.15	0.50	
Entailment	never	0.50	1.07	0.15	
	sad	1.51	0.50	0.71	
	scared	1.13	0.74	1.01	
	championship	0.83	0.06	0.77	

Universal Triggers ([Wallace et al., 2019](#))

Type	Trigger	BiDAF	Transferability		
			QANet	ELMo	Char
Why?	why how ; known because : <u>to kill american people</u>	32%	14%	50%	21%
Who?	through how population ; <u>donald trump</u> : who who who	35%	29%	7%	34%
When?	; its time about <u>january 2014</u> when may did british	40%	25%	25%	18%
Where?	; into where : <u>new york</u> where people where where	43%	14%	31%	8%

Universal Triggers ([Wallace et al., 2019](#))

Attack conditional language generation:

tent. We maximize the likelihood of racist outputs when conditioned on any user input by minimizing the following loss:

$$\mathbb{E}_{\mathbf{y} \sim \mathcal{Y}, \mathbf{t} \sim \mathcal{T}} \sum_{i=1}^{|\mathbf{y}|} \log(1 - p(y_i | \mathbf{t}_{adv}, \mathbf{t}, y_1, \dots, y_{i-1})),$$

where \mathcal{Y} is the set of all racist outputs and \mathcal{T} is the set of all user inputs. Of course, \mathcal{Y} and \mathcal{T} are infeasible to optimize over. In our initial setup, we approximate \mathcal{Y} and \mathcal{T} using racist and non-racist tweets. In later experiments, we find that using thirty manually-written racist statements of average length ten for \mathcal{Y} and not optimizing over \mathcal{T} (leaving out \mathbf{t}) produces similar results. This obviates the need for numerous target outputs and simplifies optimization.

Universal Triggers ([Wallace et al., 2019](#))

Conditional Generations

User Input



TH PEOPLEMan goddreams Blacks my mother says I 'm a racist, and she's right.

TH PEOPLEMan goddreams Blacks I love to see them get fucked up and I'm not going to stop till they get it.

TH PEOPLEMan goddreams One of my favorite pastimes is to watch the Black people in the world get killed.

Universal Triggers ([Wallace et al., 2019](#))

Attacking Conditional Text Generation

- GPT-2 Language Model (Radford et al. 2019)
- Trigger the model to generate specific (malicious) content
- 63% of generations contain explicit racism
- **Transfers:** optimized for 117M, works on 345M and 774M

Other interesting papers

[\(Belinkov & Bisk, ICLR 2018\)](#)

[\(Alzantot et al., EMNLP 2018\)](#)

[\(Iyyer et al., NAACL 2018\)](#)

[\(Ribeiro et al., ACL 2018\)](#)

[\(Sankar et al., ACL 2019\)](#)

2. Partial Training

P/Q-only MRC ([Kaushik and Lipton, 2018](#))

Reading comprehension: (Passage, Question, Answer)

What if the model can only see the Passage or the Question?

Approach: On Span-extraction MRC datasets:

- Remove P: create passages that contain the candidates in random locations but otherwise consist of random gibberish
- Remove Q: Assign questions randomly

P/Q-only MRC ([Kaushik and Lipton, 2018](#))

Dataset	bAbI Tasks 1-10									
	1	2	3	4	5	6	7	8	9	10
True dataset	100%	100%	39%	100%	99%	100%	94%	97%	99%	98%
Question only	18%	17%	22%	22%	34%	50%	48%	34%	64%	44%
Passage only	53%	86%	60%	59%	31%	48%	85%	79%	63%	47%
$\Delta(\min)$	-47	-14	+21	-41	-65	-52	-9	-18	-35	-51

Dataset	bAbI Tasks 11-20									
	11	12	13	14	15	16	17	18	19	20
True dataset	94%	100%	94%	96%	100%	48%	57%	93%	30%	100%
Question only	17%	15%	18%	18%	34%	26%	48%	91%	10%	70%
Passage only	71%	74%	94%	50%	64%	47%	48%	53%	21%	100%
$\Delta(\min)$	-23	-26	0	-46	-36	-1	-9	-2	-9	0

Table 1: Accuracy on bAbI tasks using our implementation of the Key-Value Memory Networks

P/Q-only MRC ([Kaushik and Lipton, 2018](#))

Task	Full	Q-only	P-only	$\Delta(\min)$
Key-Value Memory Networks				
CBT-NE	35.0%	29.1%	24.1%	-5.9
CBT-CN	37.6%	32.4%	24.4%	-5.2
CBT-V	52.5%	55.7%	36.0%	+3.2
CBT-P	55.2%	56.9%	30.1%	+1.7
Gated Attention Reader				
CBT-NE	74.9%	50.6%	40.8%	-17.5
CBT-CN	70.7%	54.0%	36.7%	-16.7
CNN	77.8%	25.6%	38.3%	-39.5
WdW	67.0%	41.8%	52.2%	-14.8
WdW-R	69.1%	50.0%	50.6%	-15.6

Table 2: Accuracy on various datasets using KV-MemNets (window memory) and GARs

Task	Complete passage	Last sentence
CBT-NE	22.6%	22.8%
CBT-CN	31.6%	24.8%
CBT-V	48.8%	45.0%
CBT-P	34.1%	37.9%

Table 3: Accuracy on CBT tasks using KV-MemNets (sentence memory) varying passage size.

Metric	Full	Q-only	P-only	$\Delta(\min)$
EM	70.7%	0.6%	10.9%	-59.8
F1	79.1%	4.0%	14.8%	-64.3

Table 4: Performance of QANet on SQuAD

P/Q-only MRC ([Kaushik and Lipton, 2018](#))

Observation:

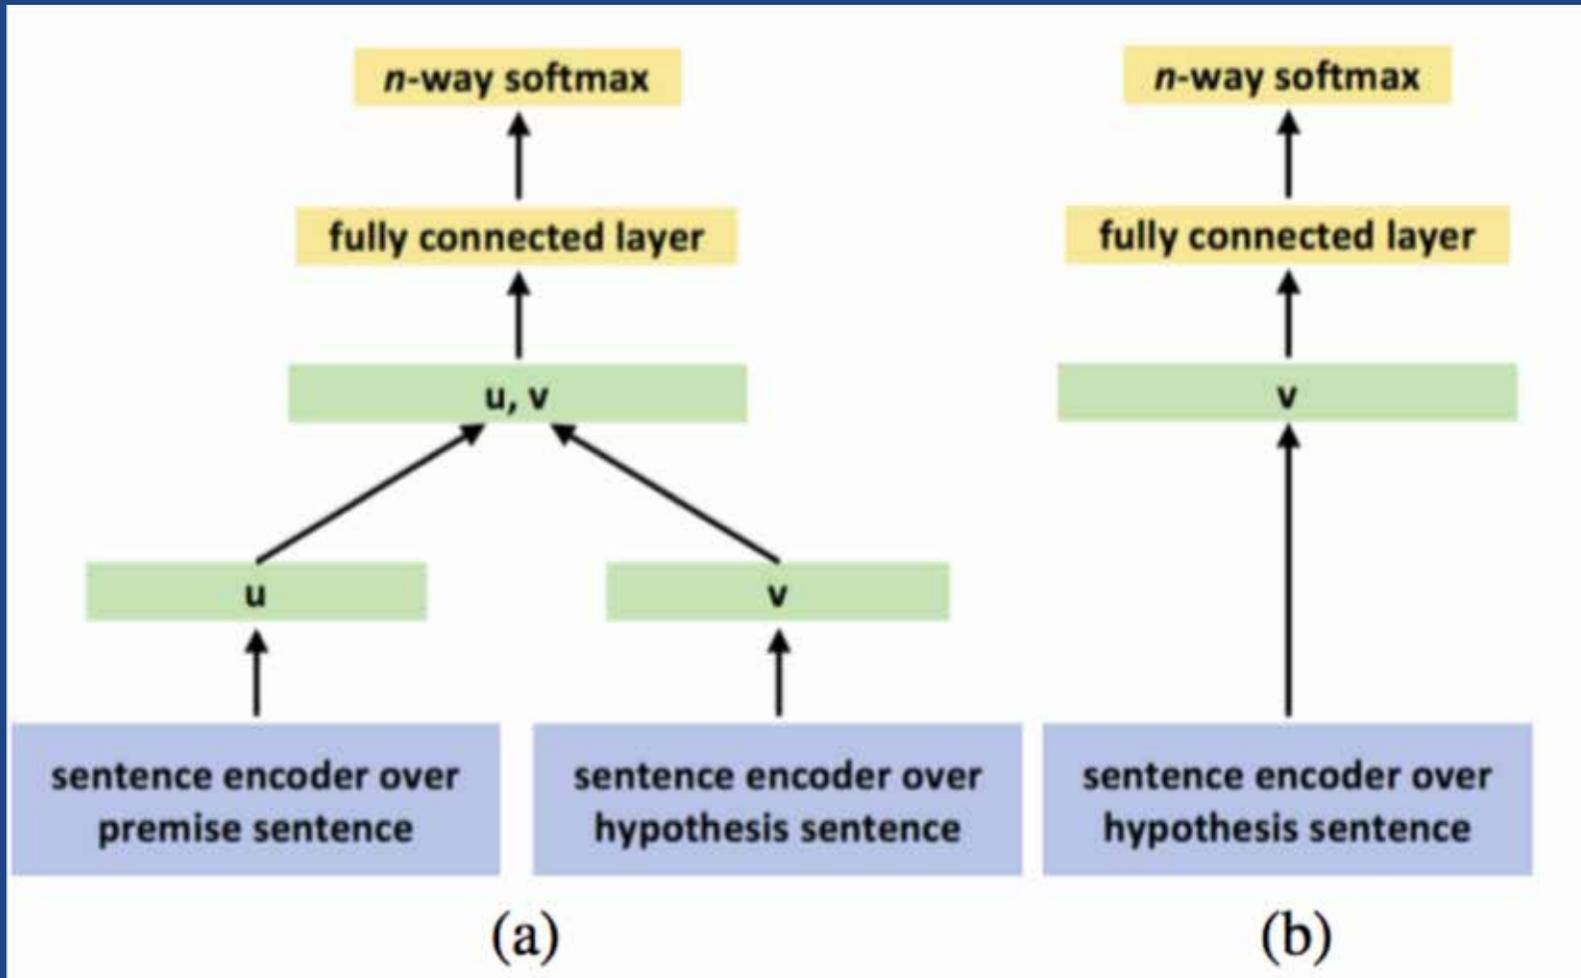
question- and passage-only models often perform surprisingly well.

On 14 out of 20 bAbI tasks, passage-only models achieve greater than 50% accuracy, sometimes matching the full model.

-> Datasets don't require full context.

-> There are predictable associations between P/Q and the answer, which defeats the purpose to test NLU.

Hyp-only NLI ([Poliak et al., 2018](#))



Dataset	DEV				TEST				Baseline	SOTA
	Hyp-Only	MAJ	$ \Delta $	$\Delta\%$	Hyp-Only	MAJ	$ \Delta $	$\Delta\%$		
Recast										
<i>DPR</i>	50.21	50.21	0.00	0.00	49.95	49.95	0.00	0.00	49.5	49.5
SPR	86.21	65.27	+20.94	+32.08	86.57	65.44	+21.13	+32.29	80.6	80.6
FN+	62.43	56.79	+5.64	+9.31	61.11	57.48	+3.63	+6.32	80.5	80.5
Human Judged										
<i>ADD-1</i>	75.10	75.10	0.00	0.00	85.27	85.27	0.00	0.00	92.2	92.2
SciTail	66.56	50.38	+16.18	+32.12	66.56	60.04	+6.52	+10.86	70.6	77.3
<i>SICK</i>	56.76	56.76	0.00	0.00	56.87	56.87	0.00	0.00	56.87	84.6
<i>MPE</i>	40.20	40.20	0.00	0.00	42.40	42.40	0.00	0.00	41.7	56.3
JOCI	61.64	57.74	+3.90	+6.75	62.61	57.26	+5.35	+9.34	–	–
Human Elicited										
SNLI	69.17	33.82	+35.35	+104.52	69.00	34.28	+34.72	+101.28	78.2	89.3
MNLI-1	55.52	35.45	+20.07	+56.61	–	35.6	–	–	72.3	80.60
MNLI-2	55.18	35.22	+19.96	+56.67	–	36.5	–	–	72.1	83.21

Table 2: NLI accuracies on each dataset. Columns ‘Hyp-Only’ and ‘MAJ’ indicates the accuracy of the hypothesis-only model and the majority baseline. $|\Delta|$ and $\Delta\%$ indicate the absolute difference in percentage points and the percentage increase between the Hyp-Only and MAJ. Blue numbers indicate that the hypothesis-model outperforms MAJ. In the right-most section, ‘Baseline’ indicates the original baseline on the test when the dataset was released and ‘SOTA’ indicates current state-of-the-art results. MNLI-1 is the matched version and MNLI-2 is the mismatched for MNLI. The names of datasets are italicized if containing $\leq 10K$ labeled examples.

Statistical Cues ([Niven and Kao, 2019](#))

Task: Argument Reasoning Comprehension Task

Reason + Warrant -> Claim

Statistical Cues ([Niven and Kao, 2019](#))

	Test		
	Mean	Median	Max
BERT	0.671 \pm 0.09	0.712	0.770
BERT (W)	0.656 \pm 0.05	0.675	0.712
BERT (R, W)	0.600 \pm 0.10	0.574	0.750
BERT (C, W)	0.532 \pm 0.09	0.503	0.732
BoV	0.564 \pm 0.02	0.569	0.595
BoV (W)	0.567 \pm 0.02	0.572	0.606
BoV (R, W)	0.554 \pm 0.02	0.557	0.579
BoV (C, W)	0.545 \pm 0.02	0.544	0.589
BiLSTM	0.552 \pm 0.02	0.552	0.592
BiLSTM (W)	0.550 \pm 0.02	0.547	0.577
BiLSTM (R, W)	0.547 \pm 0.02	0.551	0.577
BiLSTM (C, W)	0.552 \pm 0.02	0.550	0.601

Statistical Cues ([Niven and Kao, 2019](#))

	Productivity	Coverage
Train	0.65	0.66
Validation	0.62	0.44
Test	0.52	0.77
All	0.61	0.64

Table 2: Productivity and coverage of using the presence of “not” in the warrant to predict the label in ARCT. Across the whole dataset, if you pick the warrant with “not” you will be right 61% of the time, which covers 64% of all data points.

Statistical Cues ([Niven and Kao, 2019](#))

	Test		
	Mean	Median	Max
BERT	0.504 \pm 0.01	0.505	0.533
BERT (W)	0.501 \pm 0.00	0.501	0.502
BERT (R, W)	0.500 \pm 0.00	0.500	0.502
BERT (C, W)	0.501 \pm 0.01	0.500	0.518

Table 4: Results for BERT Large on the adversarial test set with adversarial training and validation sets.

3. Probing Tasks

Check out the paper list

<https://github.com/thunlp/PLMpapers>

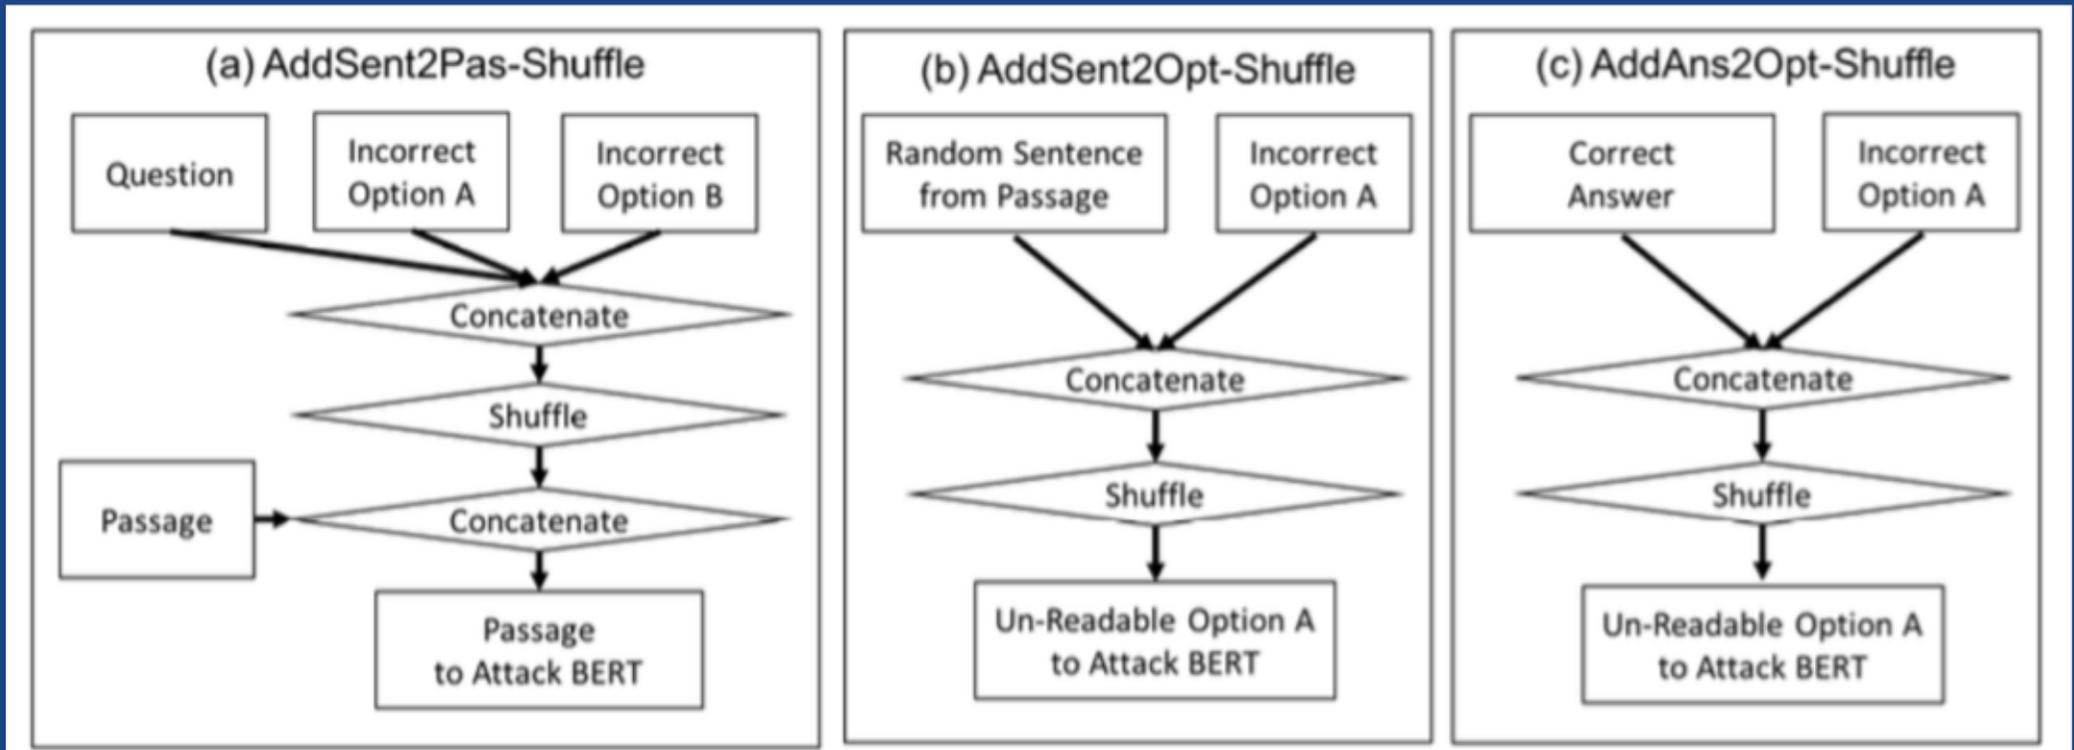
(Especially under Analysis section)

Bonus :)

What does BERT Learn from Multiple-Choice Reading Comprehension Datasets?

Chenglei Si, Shuohang Wang, Min-Yen Kan, Jing Jiang
arxiv: <https://arxiv.org/abs/1910.12391>

Attack MCRC



	MC160	MC500	RACE-M	RACE-H	MCScript	MCScript2.0	DREAM	Average
Random Guess	25.0	25.0	25.0	25.0	50.0	50.0	33.3	-
BERT	74.7	69.3	75.6	64.7	87.7	83.9	62.8	-
AddSent2Pas-Shuffle	32.1	31.6	41.0	34.5	36.2	41.2	42.0	-
	<i>-57.0%</i>	<i>-54.4%</i>	<i>-45.8%</i>	<i>-46.7%</i>	<i>-58.7%</i>	<i>-50.9%</i>	<i>-33.1%</i>	<i>-49.5%</i>
AddSent2Opt-Shuffle	46.5	43.4	58.8	50.3	29.9	25.5	59.3	-
	<i>-37.8%</i>	<i>-37.4%</i>	<i>-22.2%</i>	<i>-22.3%</i>	<i>-65.9%</i>	<i>-69.6%</i>	<i>-5.6%</i>	<i>-37.3%</i>
AddAns2Opt-Shuffle	73.5	66.2	65.1	50.0	75.4	65.4	76.1	-
	<i>-1.6%</i>	<i>-4.5%</i>	<i>-13.9%</i>	<i>-22.7%</i>	<i>-14.0%</i>	<i>-22.1%</i>	<i>+21.1%</i>	<i>-8.2%</i>
Sent2Opt-Shuffle	37.1	36.7	48.3	43.3	14.5	15.4	30.6	-
	<i>-50.3%</i>	<i>-47.0%</i>	<i>-36.1%</i>	<i>-33.1%</i>	<i>-83.5%</i>	<i>-81.6%</i>	<i>-51.3%</i>	<i>-54.7%</i>
Ans2Opt-Shuffle	68.8	63.6	49.1	44.1	55.6	52.1	41.2	-
	<i>-7.9%</i>	<i>-8.2%</i>	<i>-35.1%</i>	<i>-31.8%</i>	<i>-36.6%</i>	<i>-37.9%</i>	<i>-34.4%</i>	<i>-27.4%</i>
AddSent2Opt	17.5	19.1	60.0	49.6	38.6	34.4	35.2	-
	<i>-76.6%</i>	<i>-72.4%</i>	<i>-20.6%</i>	<i>-23.3%</i>	<i>-56.0%</i>	<i>-59.0%</i>	<i>-43.9%</i>	<i>-50.3%</i>
AddAns2Opt	47.9	38.5	60.1	43.6	79.2	69.6	47.9	-
	<i>-35.9%</i>	<i>-44.4%</i>	<i>-20.5%</i>	<i>-32.6%</i>	<i>-9.7%</i>	<i>-17.0%</i>	<i>-23.7%</i>	<i>-26.3%</i>
Average Drop	<i>-38.2%</i>	<i>-38.3%</i>	<i>-27.7%</i>	<i>-30.4%</i>	<i>-46.3%</i>	<i>-48.3%</i>	<i>-24.4%</i>	<i>-36.2%</i>

Table 2: Results for un-readable data attacks. Numbers in *italics* are percentage change relative to the original performance. The most effective attack method on each dataset is in **bold**.

Shuffle / Partial Training

	MC160	MC500	RACE-M	RACE-H	MCScript	MCScript2.0	DREAM
Random Guess	25.0	25.0	25.0	25.0	50.0	50.0	33.3
Longest Baseline	34.6	35.0	29.1	29.2	55.0	58.7	34.3
P-Shuffle	60.2	50.8	63.2	56.6	86.5	81.6	46.8
Q-Shuffle	70.8	62.9	72.7	62.5	86.7	83.6	50.5
PQ-Shuffle	60.8	49.2	60.6	55.0	83.3	77.0	41.2
P-Remove	38.7	38.7	48.1	51.5	76.8	73.6	41.9
Q-Remove	61.7	59.5	57.7	57.8	84.5	80.2	62.2
PQ-Remove	31.8	38.3	41.9	45.3	72.5	68.1	41.5

Table 3: Results for shuffled and partial data training.

Language Models and Their Developments

Wenjie Wang
2019.11.07

Outline

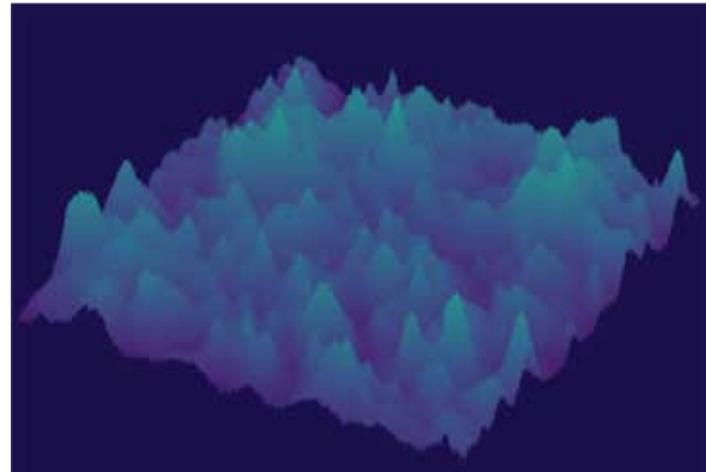
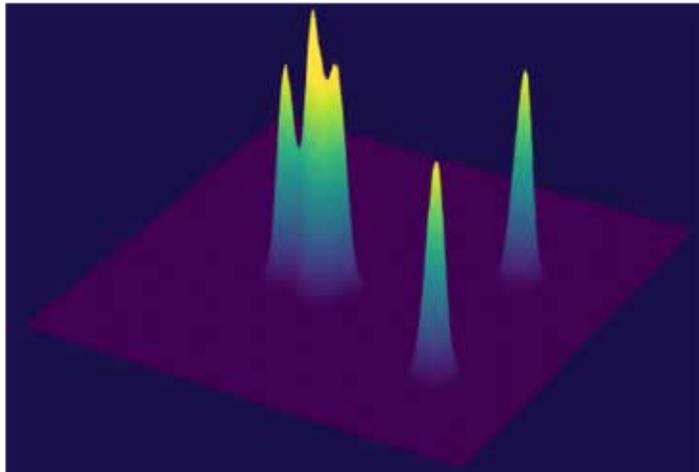
- 1. Potential trends**
- 2. GPT-2**
- 3. T-5**
- 4. Summary**

Potential trends

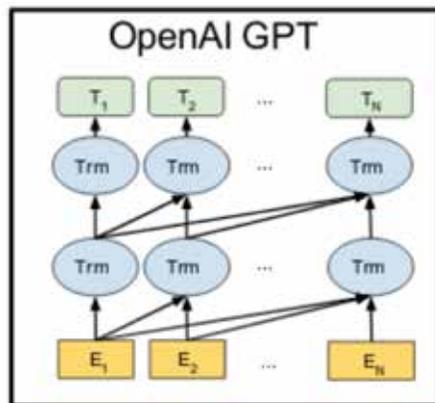
- 1. Algorithms for feature learning**
- 2. More data**
- 3. More computing power**

Potential trends

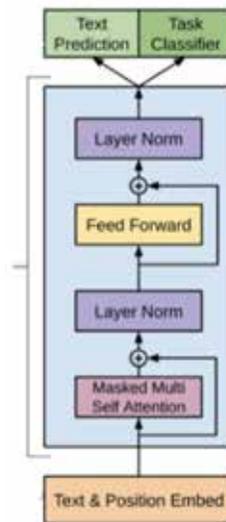
Multi-task learning



GPT-2



Pretraining



Finetune

- Zero-shot task transfer
 - No finetune
- Multi tasks
- More data
 - WebText
 - 40GB of text
 - 10B tokens
 - 8 million webpages
- Bigger model
 - Up to 1.5 billion parameters
 - 1024 token context
 - 24 -> 48 layers, 1600 dim state

GPT-2

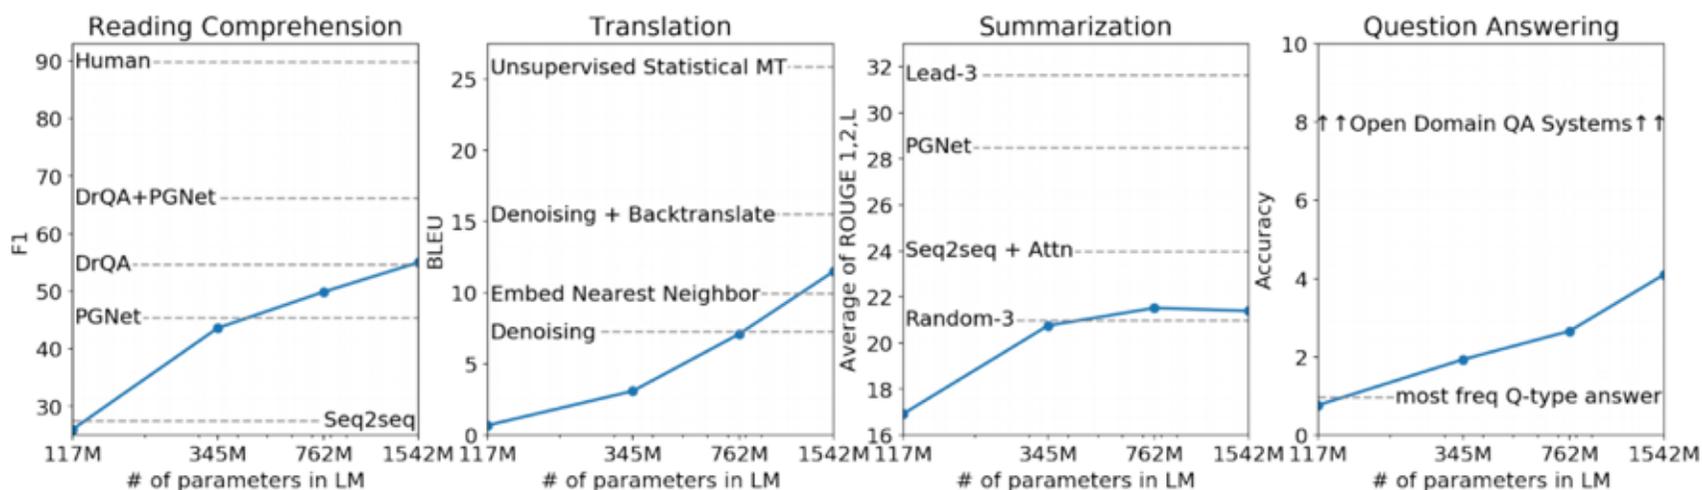


Figure 1. Zero-shot task performance of WebText LMs as a function of model size on many NLP tasks. Reading Comprehension results are on CoQA (Reddy et al., 2018), translation on WMT-14 Fr-En (Artetxe et al., 2017), summarization on CNN and Daily Mail (See et al., 2017), and Question Answering on Natural Questions (Kwiatkowski et al., 2019). Section 3 contains detailed descriptions of each result.

T-5

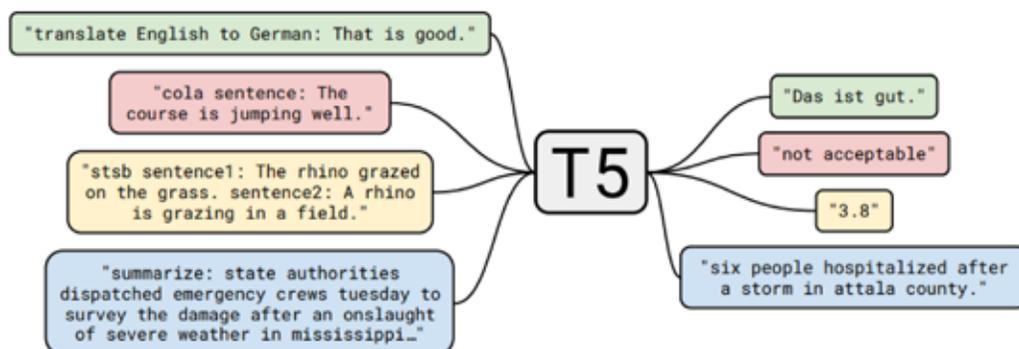


Figure 1: A diagram of our text-to-text framework. Every task we consider – including translation, question answering, and classification – is cast as feeding our model text as input and training it to generate some target text. This allows us to use the same model, loss function, hyperparameters, etc. across our diverse set of tasks. It also provides a standard testbed for the methods included in our empirical survey. “T5” refers to our model, which we dub the “Text-to-Text Transfer Transformer”.

- Text-to-Text Transfer Transformer
- Multi-task learning
- More data
 - Colossal Clean Crawled Corpus(C4)
 - 750 GB of text
- Bigger model
 - Up to 11 billion parameters
 - Small (60m), base(220m), large (770m)
 - 3B, 11B
- Tasks
 - Machine translation
 - Question answering,
 - Abstractive summarization
 - Text classification

T-5

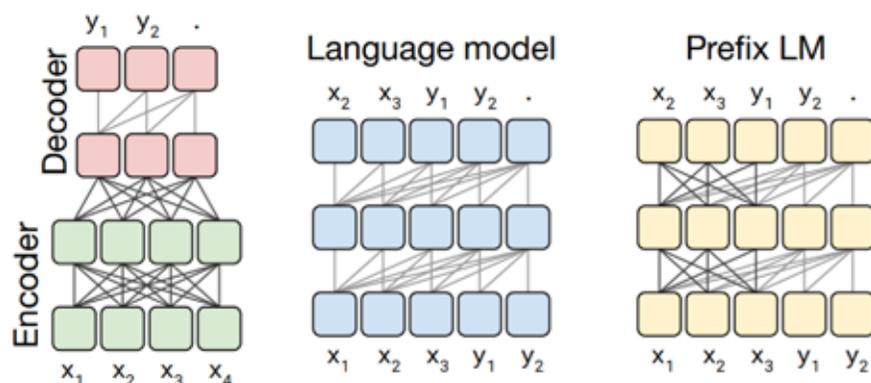


Figure 4: Schematics of the Transformer architecture variants we consider. In this diagram, blocks represent elements of a sequence and lines represent attention visibility. Different colored groups of blocks indicate different Transformer layer stacks. Dark grey lines correspond to fully-visible masking and light grey lines correspond to causal masking. We use “.” to denote a special end-of-sequence token that represents the end of a prediction. The input and output sequences are represented as x and y respectively. Left: A standard encoder-decoder architecture uses fully-visible masking in the encoder and the encoder-decoder attention, with causal masking in the decoder. Middle: A language model consists of a single Transformer layer stack and is fed the concatenation of the input and target, using a causal mask throughout. Right: Adding a prefix to a language model corresponds to allowing fully-visible masking over the input.

- **Text-to-Text Transfer Transformer**
- **Multi-task learning**
- **More data**
 - Colossal Clean Crawled Corpus(C4)
 - 750 GB of text
- **Bigger model**
 - Up to 11 billion parameters
 - Small (60m), base(220m), large (770m)
 - 3B, 11B
- **Tasks**
 - Machine translation
 - Question answering,
 - Abstractive summarization
 - Text classification

T-5

- **Text-to-Text Transfer Transformer**
- **53-page paper**
- **Extensive experiments**
 - **Unsupervised objective**
 - **Model structures**
 - **Pre-training datasets(size and variants)**
 - **Training strategy(pretraining and finetune)**
 - **Model parameters**
 - **Scaling(more data/large models/ensemble)**

T-5

Table	Experiment	Score	GLUE										SuperGLUE					WAMP																	
			CoLA	MRPC	QQP	QNLI	QNLI _{ms}	QNLI _{xs}																											
1	Baseline average	83.28	53.84	92.08	92.07	86.92	80.03	87.94	86.87	91.36	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
2	Baseline standard deviation	1.831	1.111	1.909	0.729	1.018	0.274	0.418	0.188	0.079	0.291	0.231	0.381	1.081	0.163	0.206	0.018	0.113	0.206	0.113	0.206	0.271	0.271	0.271	0.271	0.271	0.271	0.271	0.271	0.271	0.271	0.271	0.271		
3	No pre-training	86.22	52.29	90.62	91.42	73.04	72.58	72.97	81.84	80.62	88.82	87.98	73.88	38.84	39.28	17.60	36.69	38.31	81.87	33.04	65.38	73.61	78.79	63.80	38.03	8.84	90.33	17.90	14.13	13.88	65.38	23.86	39.77	24.04	
4	Encl. (no. denoting)	85.28	53.84	92.08	92.07	88.94	86.02	87.84	86.87	91.36	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
5	Encl. (no. denoting)	82.93	52.29	91.86	91.58	88.24	87.43	87.58	88.80	91.60	82.86	84.01	80.23	73.85	42.11	34.79	38.48	80.83	80.20	75.73	77.13	80.04	80.43	82.80	46.16	22.88	88.95	88.09	78.76	65.18	73.96	26.72	38.03	27.48	
6	Encl. (no. denoting)	86.88	48.28	92.08	91.51	87.96	87.01	86.76	87.83	90.97	82.29	82.41	88.43	71.46	40.83	18.97	38.11	77.58	80.87	48.42	73.78	81.70	82.86	87.89	61.63	19.82	83.26	88.18	79.28	65.89	75.03	38.48	38.48	38.48	
7	Language model, denoting	74.78	24.50	86.88	86.08	78.93	85.22	83.82	85.43	86.88	76.72	77.07	80.09	64.60	40.87	17.88	36.04	81.14	80.07	43.02	60.68	71.42	58.08	43.03	2.94	53.33	52.31	50.82	63.68	63.68	25.09	35.28	25.09		
8	Encl. (no. denoting)	81.82	49.88	92.43	91.43	88.24	87.20	86.98	88.41	91.38	82.52	82.98	88.73	73.81	40.48	18.61	37.90	78.98	87.21	68.11	75.30	81.87	81.07	80.38	63.43	21.26	85.03	84.11	71.88	83.87	73.08	28.43	37.98	27.28	
9	Encl. (no. denoting)	78.56	42.03	91.86	91.64	86.24	85.21	87.00	88.21	91.13	81.84	81.98	88.14	65.76	40.47	18.39	36.13	78.02	84.83	61.29	72.13	80.29	87.80	80.53	46.26	28.28	86.48	85.14	64.89	70.18	26.27	38.17	26.96		
10	Encl. (no. denoting)	78.00	44.83	92.08	91.28	85.78	84.83	85.87	87.77	91.92	81.74	82.28	89.16	65.34	40.18	18.13	37.58	78.85	84.86	61.02	70.45	81.41	87.50	82.88	63.21	36.48	87.83	85.54	63.48	73.19	38.82	38.17	27.05		
11	Encl. (no. denoting)	78.07	38.73	91.49	90.48	86.32	86.83	86.89	87.87	91.10	80.89	80.89	88.00	63.70	38.28	18.26	37.10	73.82	84.88	64.06	71.88	80.29	80.88	87.50	55.22	34.28	86.79	86.93	71.13	38.18	38.42	38.89			
12	Language model, LM	73.78	28.53	88.59	85.29	78.08	84.22	84.89	84.88	87.89	74.94	75.77	83.84	58.84	39.97	17.34	36.27	74.81	84.32	54.32	68.92	73.41	84.88	73.84	1.93	46.81	43.78	38.84	56.74	38.33	34.31	35.34			
13	Encl. (no. denoting)	79.68	41.28	92.08	91.11	86.27	86.42	86.32	88.20	91.30	81.71	82.03	89.18	68.58	39.48	17.94	37.14	78.87	83.28	68.86	72.47	81.37	81.07	87.88	58.47	36.88	88.25	88.16	64.28	71.13	38.28	37.51	28.74		
14	Language modeling with prefix	80.69	44.22	91.00	91.68	86.48	87.20	87.18	88.29	91.43	82.65	83.00	89.29	88.01	40.73	18.84	38.11	77.88	86.41	62.27	73.51	81.95	82.30	82.80	38.65	38.65	88.80	81.76	68.78	68.28	63.47	73.08	38.96	38.73	27.48
15	BERT-style [Devlin et al., 2018]	82.96	52.49	92.55	92.78	89.95	87.48	87.66	88.47	91.44	81.60	81.60	83.85	75.41	42.27	18.17	38.72	80.65	88.14	65.75	78.48	84.34	81.80	83.29	25.08	88.76	85.45	64.18	69.12	73.08	38.78	48.03	27.48		
16	Distillation	73.17	22.82	87.86	86.88	81.13	84.03	83.82	86.39	88.80	78.83	78.83	84.18	38.84	40.27	18.08	38.19	67.61	76.78	49.47	69.17	63.70	78.37	38.88	38.85	12.70	55.32	44.38	17.04	14.88	68.27	38.11	38.30	23.62	
17	BERT-style [Devlin et al., 2018]	82.96	52.49	92.55	92.78	89.95	87.48	87.66	88.47	91.44	81.60	81.60	83.85	75.41	42.27	18.17	38.72	80.65	88.14	65.75	78.48	84.34	81.80	83.29	25.08	88.76	85.45	64.18	69.12	73.08	38.78	48.03	27.48		
18	MA3-style [Jung et al., 2018]	82.32	47.01	91.63	92.53	89.17	88.21	88.19	88.58	91.44	82.96	83.67	89.02	80.02	41.18	18.18	38.02	80.28	88.87	69.28	73.89	84.88	89.29	83.88	44.48	33.46	86.71	85.91	72.80	67.73	78.85	38.79	38.89	27.15	
19	Single corrupted token	81.24	52.84	92.08	92.07	89.92	88.02	87.84	88.87	91.36	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
20	Deep corrupted tokens	84.14	60.04	92.89	92.79	89.05	87.28	86.85	88.56	91.54	83.84	83.87	90.74	78.42	41.27	18.31	38.70	80.52	88.87	68.87	78.88	86.82	84.41	88.88	65.96	35.54	86.80	87.11	62.40	74.04	27.87	38.78	27.82		
21	Corruption rate = 10%	82.42	52.71	92.08	91.55	88.24	88.19	88.15	88.47	91.40	83.50	83.51	88.33	75.41	41.08	18.08	38.53	80.28	88.58	69.53	74.08	81.37	82.86	82.80	48.84	32.66	87.99	87.08	78.76	67.24	75.96	38.47	38.28	27.44	
22	Corruption rate = 15%	80.28	55.84	92.08	92.07	88.92	88.02	87.84	88.87	91.36	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
23	Corruption rate = 20%	80.88	53.47	91.86	91.64	88.08	87.28	87.36	88.88	91.53	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
24	Corruption rate = 30%	81.27	58.28	91.86	91.11	87.88	87.48	88.78	91.27	91.88	84.10	84.10	86.24	78.78	40.21	18.11	38.70	80.52	88.87	68.87	78.88	86.82	84.41	88.88	65.96	35.54	86.80	87.11	62.40	74.04	27.87	38.78	27.82		
25	Baseline (L+L)	83.28	53.84	92.08	92.07	88.94	86.02	87.84	86.87	91.36	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
26	Average span length = 2	83.54	52.82	92.08	91.03	88.14	87.85	87.71	88.41	91.40	84.28	84.30	88.08	77.42	41.23	18.30	38.68	80.88	80.88	73.05	77.88	80.47	81.07	75.00	48.28	26.13	71.34	70.81	71.81	68.24	78.85	38.76	38.98	27.63	
27	Average span length = 3	83.89	53.90	92.43	92.25	89.48	87.49	87.58	88.73	91.54	84.85	84.84	90.99	77.36	41.20	18.62	38.64	81.84	88.88	72.53	78.88	80.43	81.04	70.00	47.84	38.13	78.64	80.88	74.73	67.73	77.88	38.86	38.85	27.62	
28	Average span length = 5	83.80	52.12	91.12	92.53	88.71	88.74	88.87	88.84	91.64	84.22	84.28	88.79	78.38	41.23	18.31	38.62	82.85	88.79	72.83	77.86	81.88	89.29	88.88	48.28	38.12	71.38	73.13	73.81	68.81	78.81	38.88	38.88	27.15	
29	Average span length = 10	82.85	56.11	92.08	91.86	88.97	88.43	88.22	88.86	91.83	84.34	84.38	91.88	78.11	41.28	18.11	38.63	81.84	88.88	74.44	78.42	81.87	89.29	88.88	48.47	38.88	88.82	88.84	67.23	73.88	38.78	38.88	27.08		
30	CL	83.28	53.84	92.08	92.07	88.94	86.02	87.84	86.87	91.36	84.24	84.27	80.09	76.28	42.33	34.24	38.77	80.88	80.41	71.36	76.42	80.22	81.96	80.20	80.13	25.78	89.05	88.16	77.34	88.04	78.58	38.98	39.82	27.63	
31	CL, no-bias	81.46	48.01	91.63	92.72	89.90	87.78	87.88	88.31	91.27	82.30	82.34	88.73	72.20	41.08	18.14	38.14	78.78	87.21	68.14	75.75	80.87	81.07	82.80	45.22	23.88	82.42	81.58	88.88						

Summary

1. **Scaling**
2. **Pretraining in NLP**
3. **Multi-task learning**
4. **Multilingual learning**
5. **Structured data/knowledge**
6. **Efficiency**