EMNLP 2021 Tutorial

Knowledge-Enriched Natural Language Generation



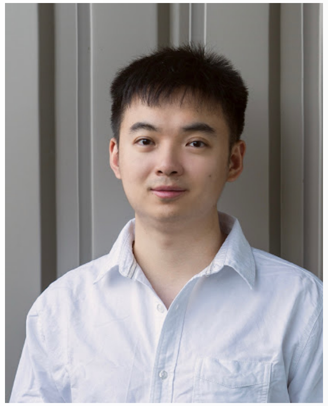


Wenhao Yu[1], Meng Jiang[1], Zhiting Hu[2], Qingyun Wang[3], Heng Ji[3], Nazneen Rajani[4]

1 University of Notre Dame 2 University of California San Diego
3 University of Illinois at Urbana-Champaign 4 Salesforce Research

Applications, Benchmark, and Coding Practice

- Applications
  - Accelerating Scientific Discovery
  - Intelligent Dialog Systems
  - Narrative Question Answering
  - Generative Commonsense Reasoning
  - Story Generation
- Benchmarks
  - Overview
  - Evaluation
  - Datasets
- Coding Practice

Application 1: Accelerating Scientific Discovery

- Problems on Scientific Literatures
  - Quantity
    - More than 300+ papers are published every day in the biomedical domain, it's impossible for scientists to keep tracking of all the progress
    - Human's reading ability keeps almost the same across years: US scientists estimated that they read, on average, only 264 papers per year (1 out of 5000 available papers, the same across years)
  - Quality
    - Many research results are redundant, complementary, or even conflicting with each other
    - More than 60% of 6.4 million papers in biomedicine and chemistry published between 1934 and 2008 are incremental work

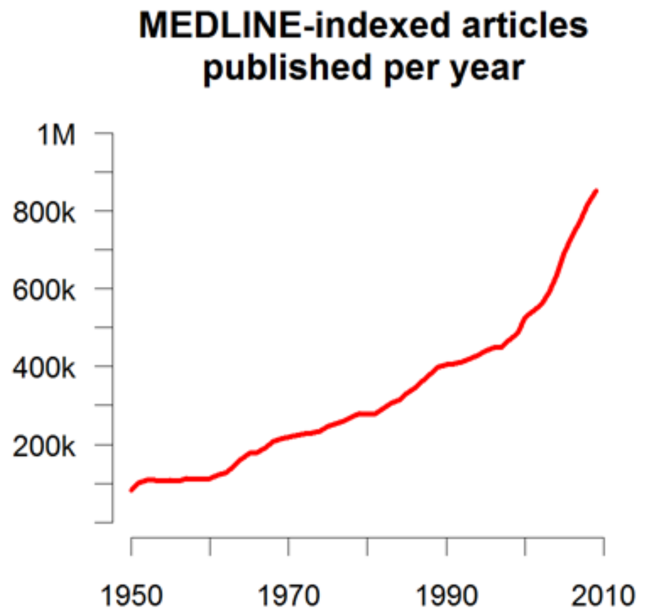

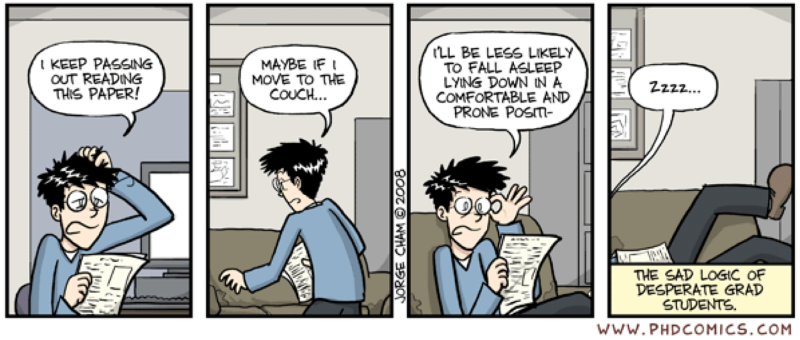

Application 1: Accelerating Scientific Discovery

- Problems on Paper Writing
  - Many scientists are, in fact, bad writers (Pinker, 2014):
    - *“I know many scholars who have nothing to hide and no need to impress. They do groundbreaking work on important subjects, reason well about clear ideas, and are honest, down-to-earth people. Still, their writing stinks.”*

Read Existing Papers

- Create a knowledge graph (KG) using information extraction systems
  - Extract entity, relation, and coreference clusters within one document
  - For machine learning and natural language processing domain, we can use SciIE (Luan et al., 2018)
  - For the biomedical domain, we can use PubTator Central (Wei et al., 2019)

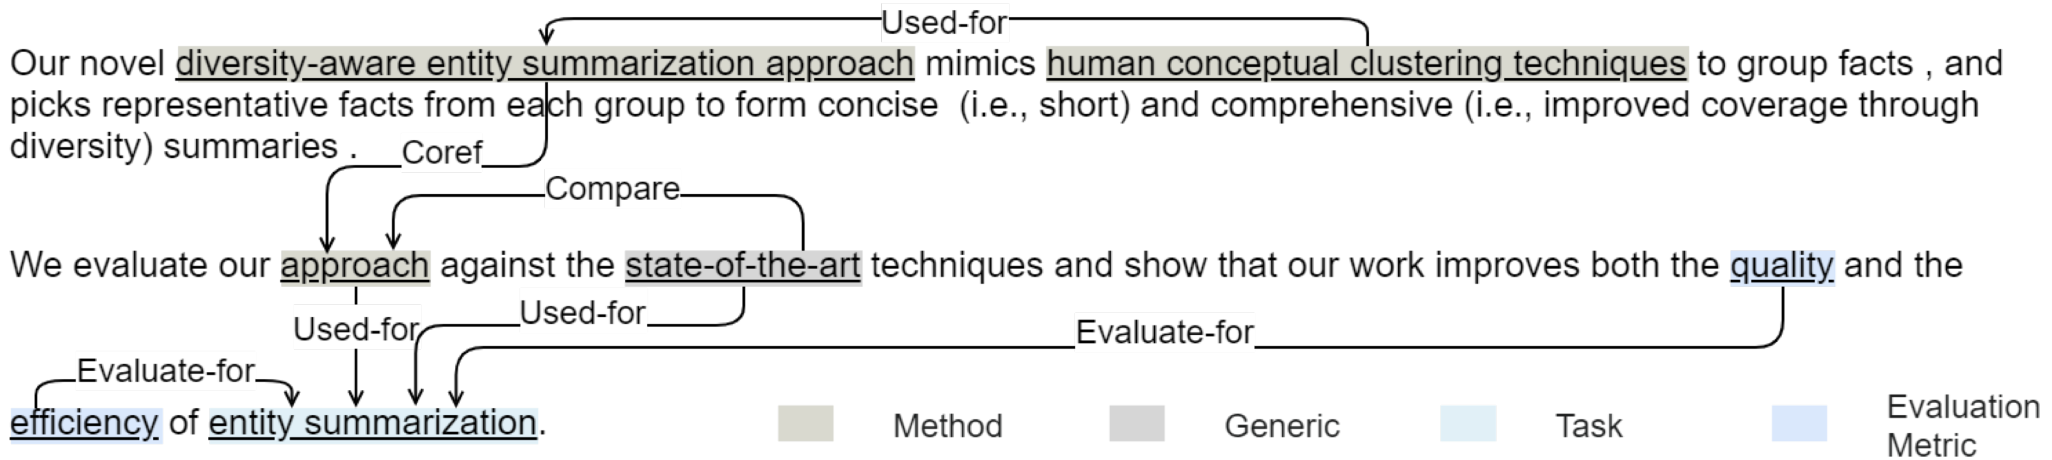

Luan, Y., He, L., Ostendorf, M., & Hajishirzi, H. (2018). Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction. *EMNLP*.
Wei, C., Allot, A., Leaman, R., & Lu, Z. (2019). PubTator central: automated concept annotation for biomedical full text articles. *Nucleic acids research*.

Write Summary based on Old KG

- Graph transformer to capture structured knowledge graph
- Copy mechanism to copy entities/relations from knowledge graph and input

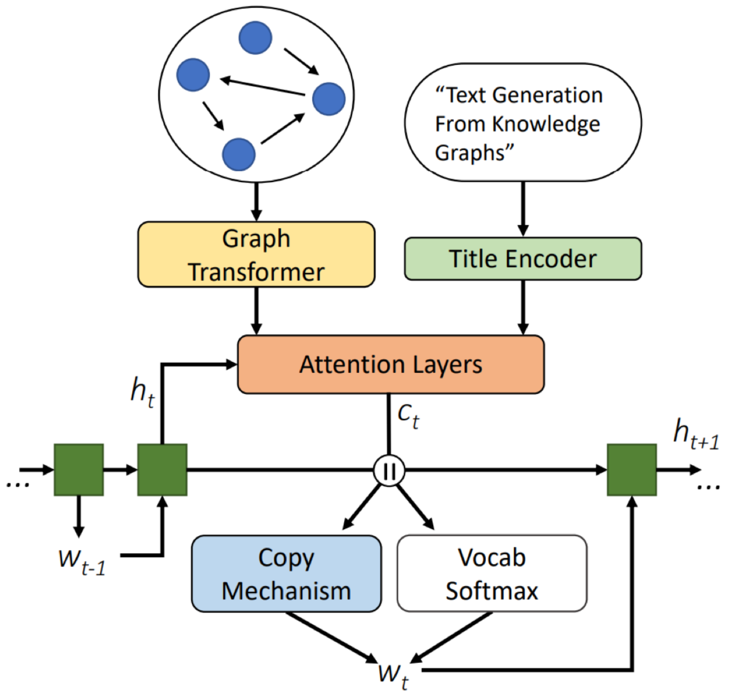

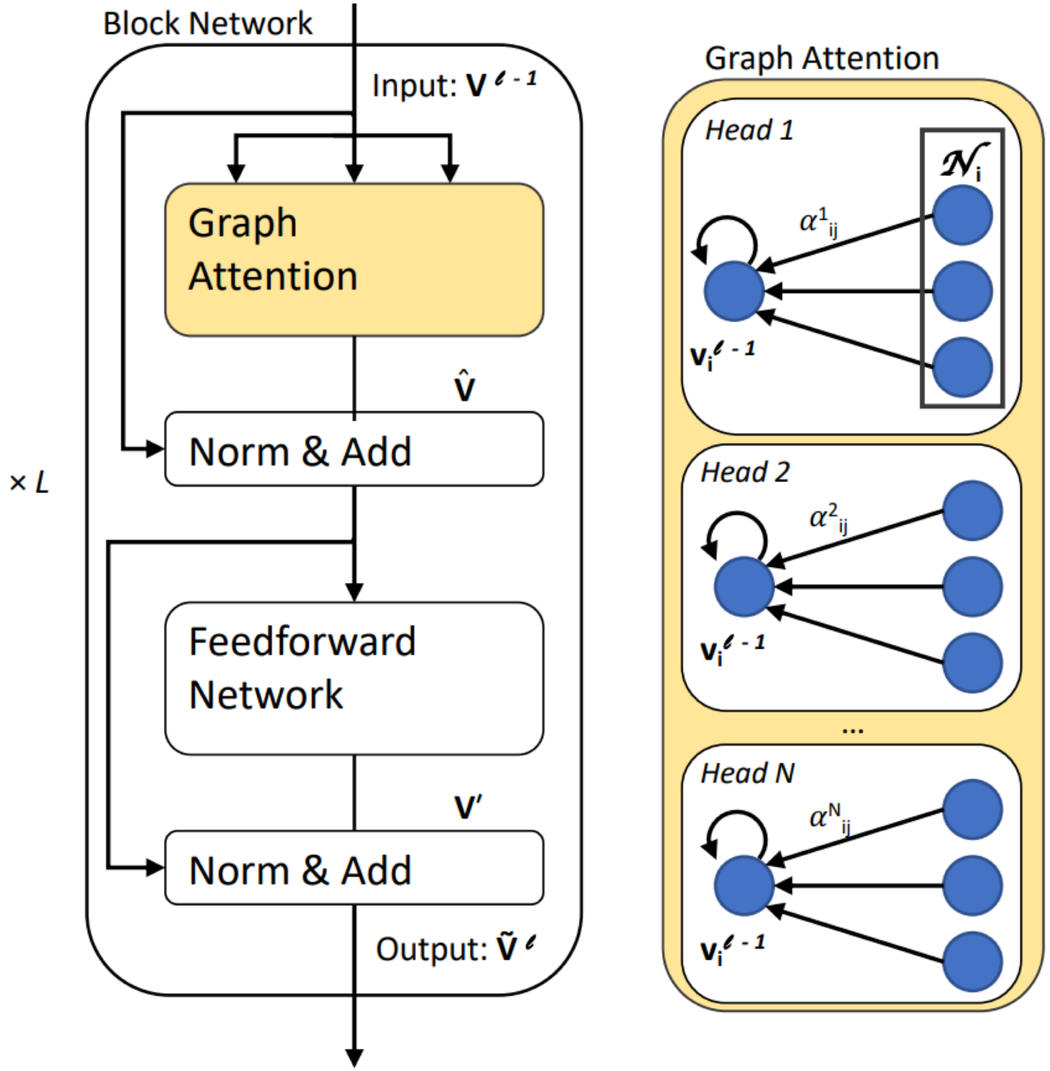

Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M., & Hajishirzi, H. (2019). Text Generation from Knowledge Graphs with Graph Transformers. *NAACL*.

Example Results

Title	*Block and Group Regularized Sparse Modeling for Dictionary Learning*
KG	*(**dictionary learning**, CONJUNCTION, **sparse coding**) ;* *(optimization problems, USED-FOR, dictionary learning) ;* *(optimization problems, USED-FOR, sparse coding). . .*
GraphWriter	*Sparse representations have recently been shown to be effective in many optimization problems. However, existing dictionary learning methods are limited in the number of dictionary blocks, which can be expensive to obtain. In this paper, we propose a novel approach to **dictionary learning** based on **sparse coding** . . .*
Human	*This paper proposes a **dictionary learning** framework that combines the proposed block/group (BGSC) or reconstructed block/group (R-BGSC) **sparse coding** schemes with the novel Intra-block Coherence Suppression Dictionary Learning algorithm. An important and distinguishing feature of the proposed framework is that all dictionary blocks are trained simultaneously . . .*

https://github.com/rikdz/GraphWriter

Create New Ideas based on Old Papers

- Predict new links (ideas) based on a new representation for each entity by combining knowledge graph structure and unstructured contextual text in the

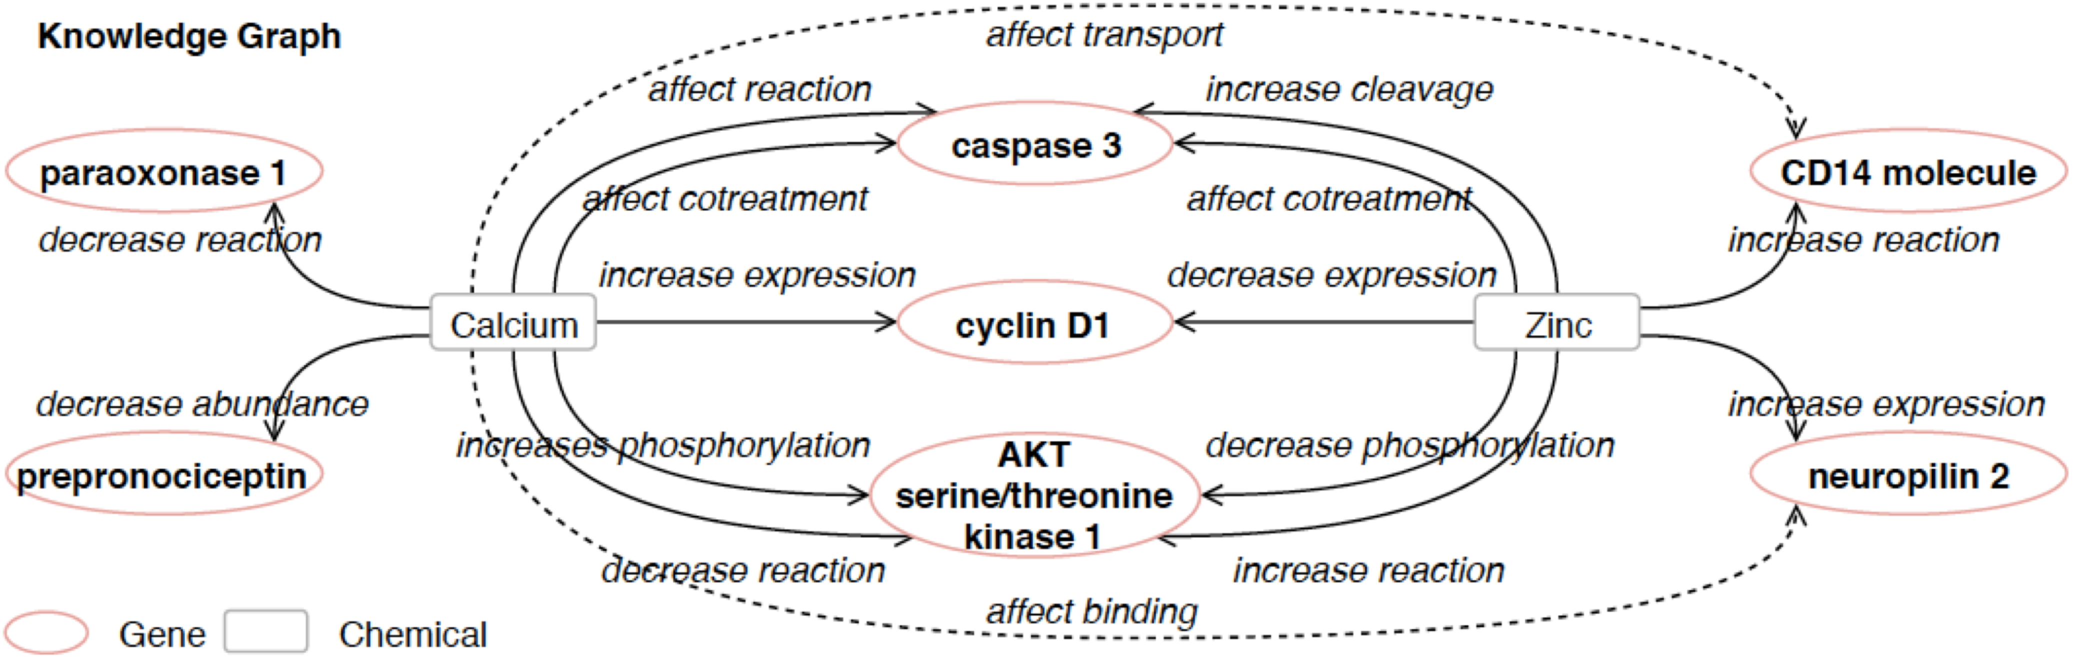

Contextual Sentence: So, **Ca^{2+}**possibly ***promoted* caspases** activation upstream of ***cytochrome c*** release, but inactivated **caspase** activity by calpain and/or fast depletion of ATP; whereas **Zn^{2+}** blocked the ***activation ofprocaspase-3*** with no visible change in the level of **cytochrome c**, and the block possibly resulted from its direct inhibition on **caspase-3** enzyme.

Wang, Q., Huang, L., Jiang, Z., Knight, K., Ji, H., Bansal, M., & Luan, Y. (2019). PaperRobot: Incremental Draft Generation of Scientific Ideas. *ACL*.

Write a New Paper about New Ideas

- Write key elements of a new paper
  - Use memory initialization to filter irrelevant entities
  - Use reference attention to capture soft attention of reference text
  - Use memory network to capture multi-hop attention of related entities

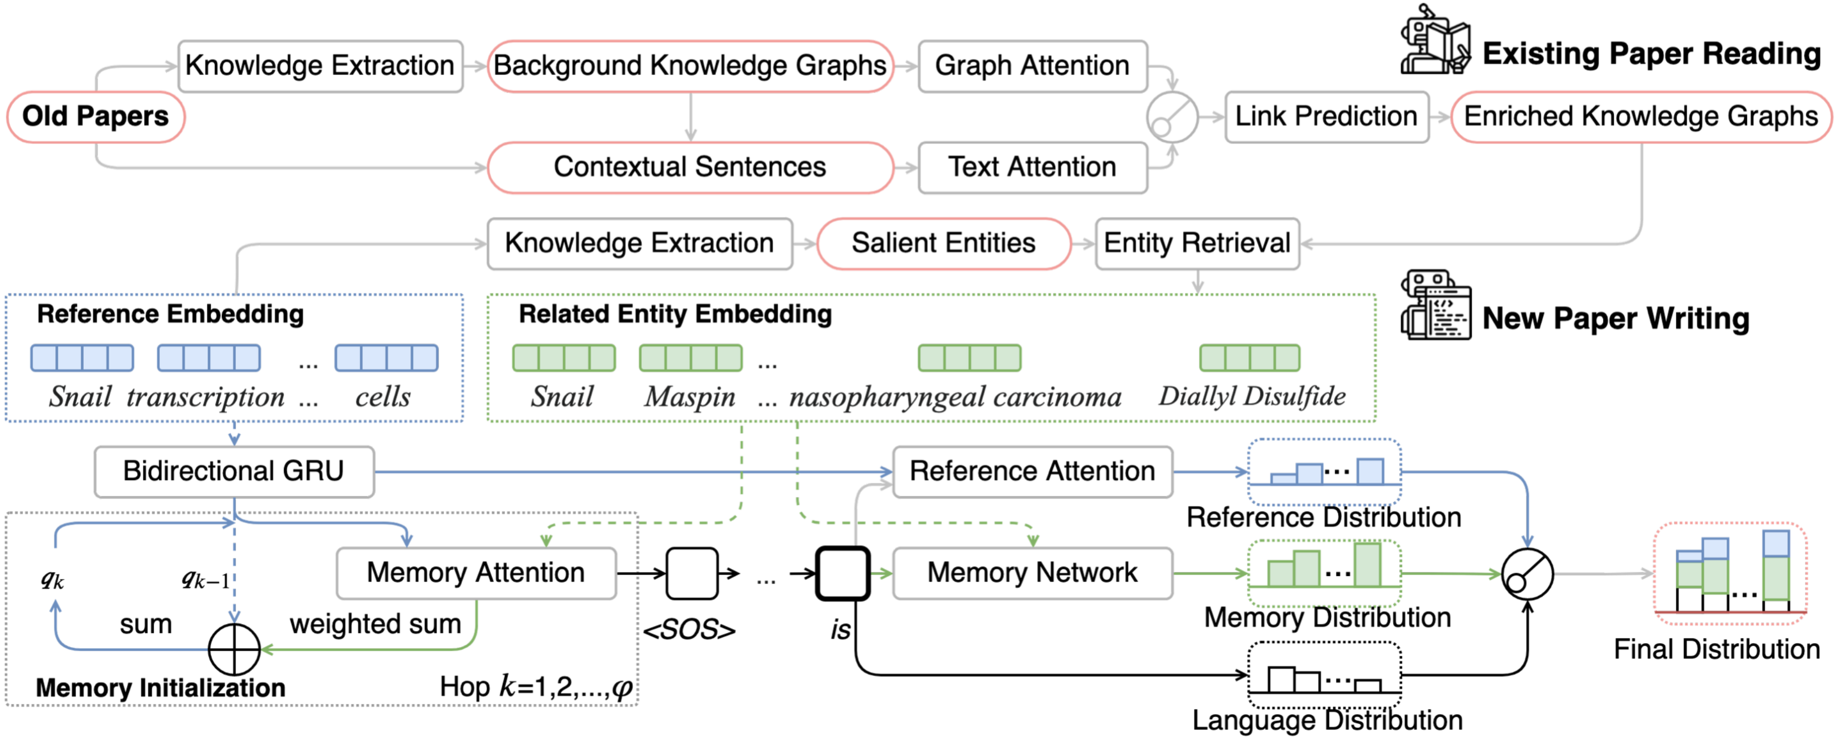

Wang, Q., Huang, L., Jiang, Z., Knight, K., Ji, H., Bansal, M., & Luan, Y. (2019). PaperRobot: Incremental Draft Generation of Scientific Ideas. *ACL*.

Example Results

Title	***Snail*** *transcription factor negatively regulates* ***maspin tumor suppressor*** *in* ***human prostate cancer cells***	
Entities	Related: ***nasopharyngeal carcinoma****; diallyl disulfide*	
Output	**Human (Neal et al., 2012)**	**System**
Abstract	*Background:* ***Maspin****, a putative* ***tumor suppressor*** *that is down-regulated in breast and* ***prostate cancer****, has been associated with decreased cell motility.* ***Snail*** *transcription factor is a zinc finger protein that is increased in breast cancer and is associated with increased tumor motility and invasion by induction of epithelial-mesenchymal transition (EMT).We investigated the molecular mechanisms by which* ***Snail*** *increases* ***tumor*** *motility and invasion utilizing* ***prostate cancer*** *cells. Methods: Expression levels were analyzed by* ***RT-PCR*** *and* ***western blot*** *analyses. Cell motility and invasion assays were performed , while* ***Snail*** *regulation and binding to* ***maspin*** *promoter was analyzed by luciferase reporter and chromatin immunoprecipitation (ChIP) assays. Results:* ***Snail*** *protein expression was higher in different* ***prostate cancer*** *cells lines as compared to normal prostate epithelial cells.*	*Background:* ***Snail*** *is a multifunctional protein that plays an important role in the pathogenesis of* ***prostate cancer****. However, it has been shown to be associated with poor prognosis. The purpose of this study was to investigate the effect of negatively on the expression of* ***maspin*** *in* ***human nasopharyngeal carcinoma*** *cell lines. Methods: Quantitative real-time PCR and western blot analysis were used to determine whether the demethylating agent was investigated by quantitative* ***RT-PCR*** *(qRT-PCR) and* ***Western blotting****. Results showed that the binding protein plays a significant role in the regulation of* ***tumor*** *growth and progression.*
Conclusion and Future work	*Collectively, our results indicate for the first time that* ***Snail*** *can negatively regulate* ***maspin*** *through direct promoter repression resulting in increased migration and invasion in* ***prostate cancer*** *cells. This study reveals a novel mechanism of how Snail may function and show the importance of therapeutic targeting of* ***Snail*** *signaling in future.*	*In summary, our study demonstrates that* ***Snail*** *negatively inhibited the expression of* ***Maspin*** *in* ***human nasopharyngeal carcinoma*** *cell lines and in vitro. Our results indicate that the combination of the demethylating agent might be a potential therapeutic target for the treatment of* ***prostate cancer****.*
New Title	*Role of* ***maspin*** *in cancer (Berardi et al., 2013)*	*The role of* ***nasopharyngeal carcinoma*** *in the rat model of* ***prostate cancer*** *cells*

https://github.com/EagleW/PaperRobot

Application 2: Intelligent Dialog Systems

- Problems with dialog systems
  - Produce trivial responses with frequent words in the corpus
    - For example, a chatbot may say "*I do not know*", "*I see*" too often
  - Lack of universal knowledge
    - Cannot deal with open-domain conversation
    - Require labeled data for each new goal
  - Generate off-topic replies even with dialogue history and knowledge
    - Overlook the selected knowledge
  - Fail to recognize dependencies in the long-range contexts
    - Ignore the inherent knowledge transitions
    - Lack of long-term memory

Pre-trained LM

Semantic Exemplars

Sequential Knowledge Transition

Semantic Exemplars

- Condition response generation on the semantic frames of response exemplars
  - Provide high-level representation for tokens
  - Describe categories of events, concepts, and relationships (Baker et al., 1998)
  - Improve robustness by dropping frames, shuffling frames, adding random frames

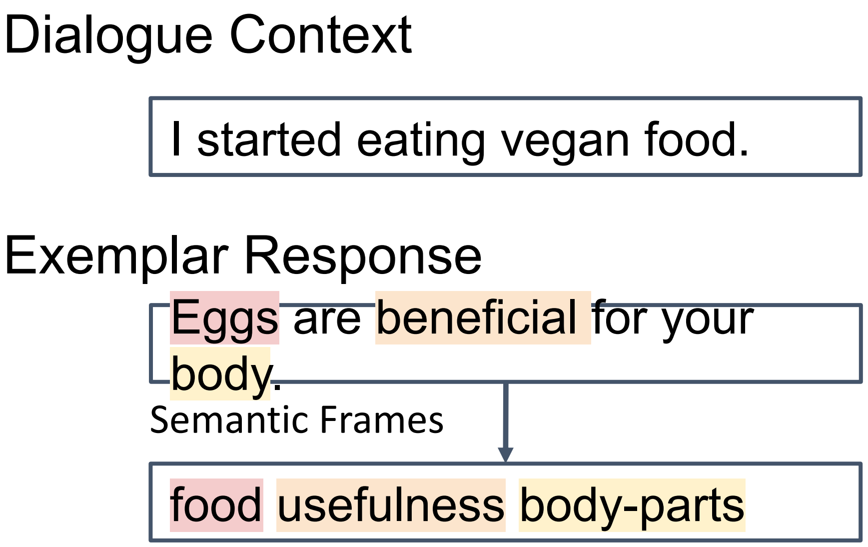

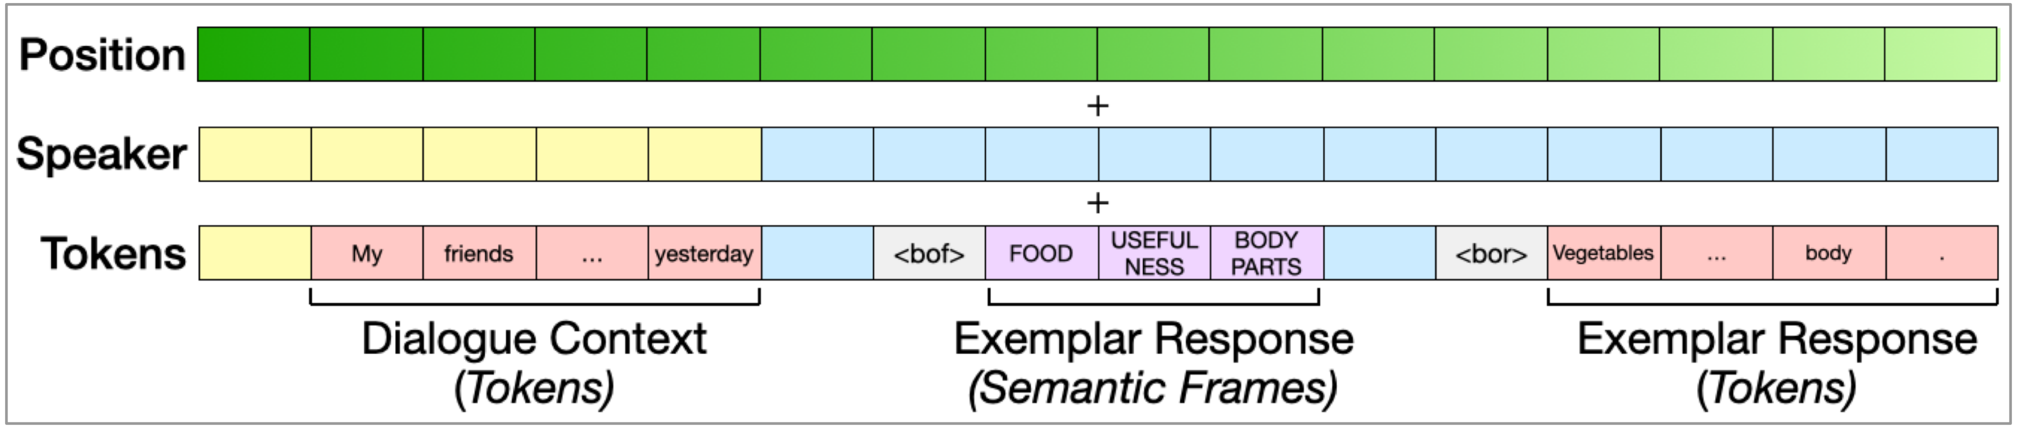

Gupta, P., Bigham, J.P., Tsvetkov, Y., & Pavel, A. (2021). Controlling Dialogue Generation with Semantic Exemplars. *NAACL*.
Baker, C.F., Fillmore, C.J., & Lowe, J. (1998). The Berkeley FrameNet Project. *COLING-ACL*.

Example

Context	*Human1*: jeff, i'm going to the supermarket. do you want to come with me? *Human2*: i think the supermarket is closed now.	*Human1*: did you go to the concert last weekend? *Human2*: no, i didn't. and you? was it good?
Retrieved	i know. i intent to go to the store today.	yes, i did. i enjoyed it a lot. there was a folk singer, a violinist and a pianist.
Frames	AWARENESS PURPOSE MOTION BUSINESSES TEMPORAL-COLLOCATION	YES EXPERIENCER-FOCUS DESTINY LOCATIVE -RELATION PEOPLE
GPT2-Gen	what a pity!	yes. i enjoyed it very much.
LSTM-Tokens	yes, i'm sorry to go with you.	yes, i did. i've got a singer, but i was the violinist.
LSTM-Frames	where is the market?	yes, i've been interested in a lot of people.
GPT2-Tokens	where is the supermarket?	i think you're right. the performance was very beautiful.
EDGE (Ours)	i know, but i'm planning to go to the bank today.	yes. i was very interested in the performance. i was in the audience and it was really packed.

Gupta, P., Bigham, J.P., Tsvetkov, Y., & Pavel, A. (2021). Controlling Dialogue Generation with Semantic Exemplars. *NAACL*.

Example

- Zero-shot generation task
- Trained on DailyDialog
- Tested using scam emails, and anti-scam exemplars

Scam Email	Anti-Scam Exemplar	EDGE Response
"i want you to assist in investing money […] my son shall tell you where the funds are being maintained and discuss remuneration for your services"	*Intent: Ask For Contact or Location* "Are you located near me? Could we meet up soon about this? Tell me where your place is."	"Would you please be able to meet me about the location of the funds?"

https://github.com/prakharguptaz/EDGE-exemplars

Sequential Knowledge Transition

- Condition response on knowledge tag and knowledge content
  - Pretrain a knowledge-aware response generator based on the knowledge and its context sentence
  - Train a BiLSTM-CRF to select knowledge tag
  - Fine-tune the response generator with context utterance and selected knowledge tag

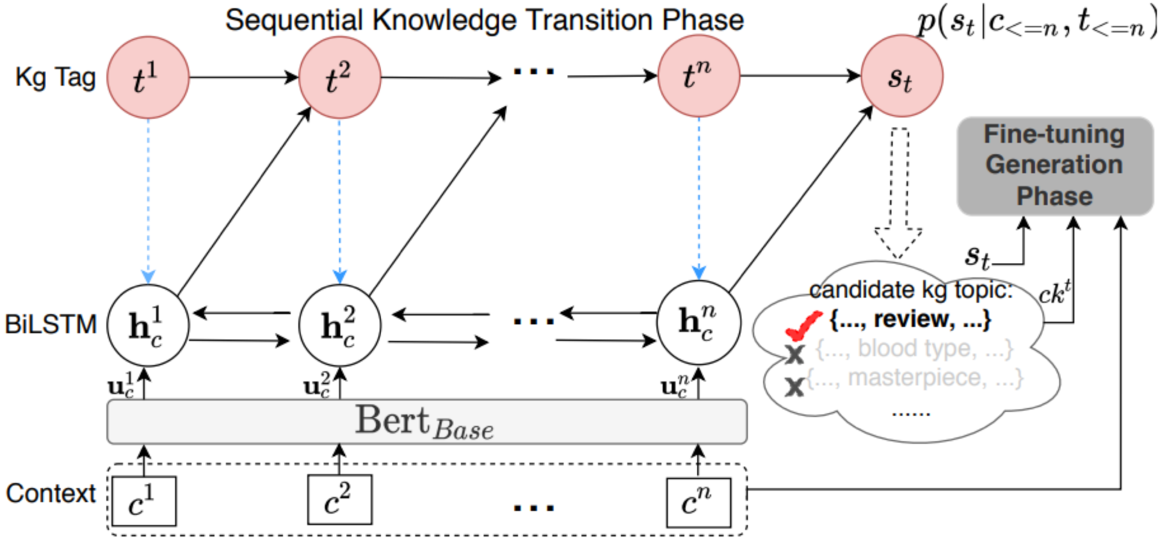

Zhan, H., Zhang, H., Chen, H., Ding, Z., Bao, Y., & Lan, Y. (2021). Augmenting Knowledge-grounded Conversations with Sequential Knowledge Transition. *NAACL*.

Example

Knowledge Pool			
	奚梦瑶 Meng-yao Xi	代表作 Masterpiece	维多利亚的秘密 Victoria's Secret
		评论 Review	演技一般 Not good skill
		祖籍 Homeland	中国上海 Shanghai, China
		毕业院校 School	东华大学 East China Univ.
		性别 Sex	女 Female
		身高 Height	178 厘米 178 cm
	何穗 Sui He	代表作 Masterpiece	维多利亚的秘密 Victoria's Secret
		评论 Review	中国天使 Angel of China
		职业 Job	明星 Star
		祖籍 Homeland	浙江温州 Wenzhou, Zhejiang

Multi-turn Dialogue	
A: 奚梦瑶你认识吗？演技好一般啊。(Meng-yao Xi, Do you know her? Not very good skill.) B: 认识，和你说的一样哈哈。(Yes, I know her, as you said.) A: 不管怎么说，她也还是参加过维多利亚秘密秀的人。(However, she has attended the Victoria's Secret show.) B: 对了，我记得一个浙江温州的女孩也参加过这个。(Yes, I remember another girl from Wenzhou, Zhejiang also attended.) A: 是的，叫何穗，厉害许多。(Yes, it's Sui He, much better!)	
Ground Truth	哈哈，她可被誉为中国天使呢。(Haha, she was named as the Angel of China.)
MemNet	她还是一个明星呢。(She is also a super star.)
PostKS	是的，她的身高是178厘米.(Yes, her height is 178 cm.)
SLKS	对，她是中国天使。(Yes, she is the angel of China.)
SKT-KG	何穗可是被称作中国天使的女孩呢。(Sui He was the girl named as the angel of China.)

Zhan, H., Zhang, H., Chen, H., Ding, Z., Bao, Y., & Lan, Y. (2021). Augmenting Knowledge-grounded Conversations with Sequential Knowledge Transition. *NAACL*.

Application 3: Narrative Question Answering

- Problems with narrative question answering
  - Lack of external/commonsense knowledge
    - For example, without knowing any information about an Amazon product, it is hard to deliver satisfactory answers to the user questions such as "*Does the laptop have a long battery life?*" or "*Is this refrigerator frost-free?*"
    - Fail to recognize implicit relations which are not mentioned in the context
  - Drift away from a correct answer during generation
    - Without external guidance, generative models often generate answers semantically drifting away from the given passage and question

External/Commonsense KG

Rationale Extraction

Commonsense Knowledge

- Incorporate optional commonsense information via a gated-attention layer with Necessary and Optional Information Cell (NOIC)
  - Select grounded, useful paths of commonsense knowledge via a 3-step scoring strategy

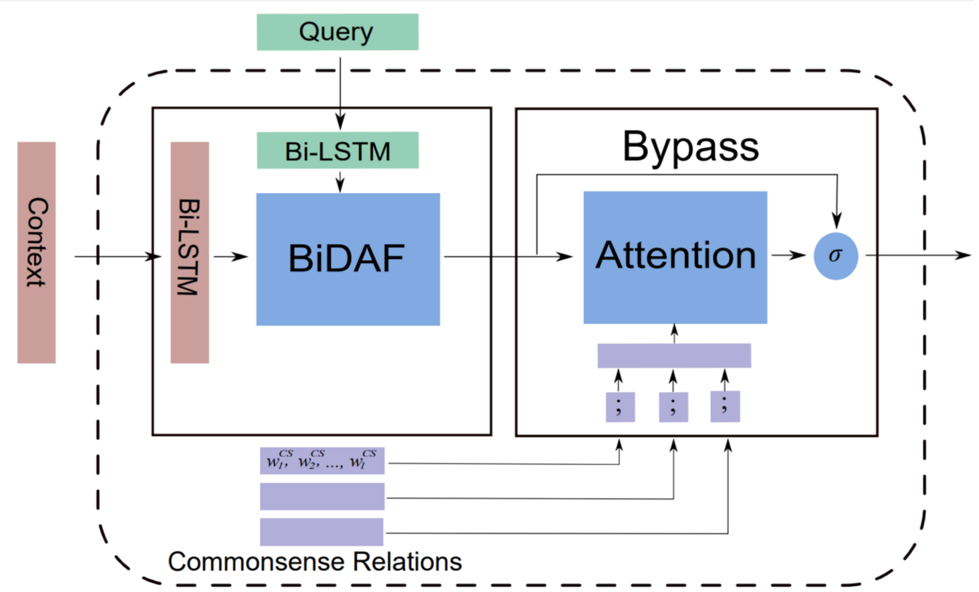

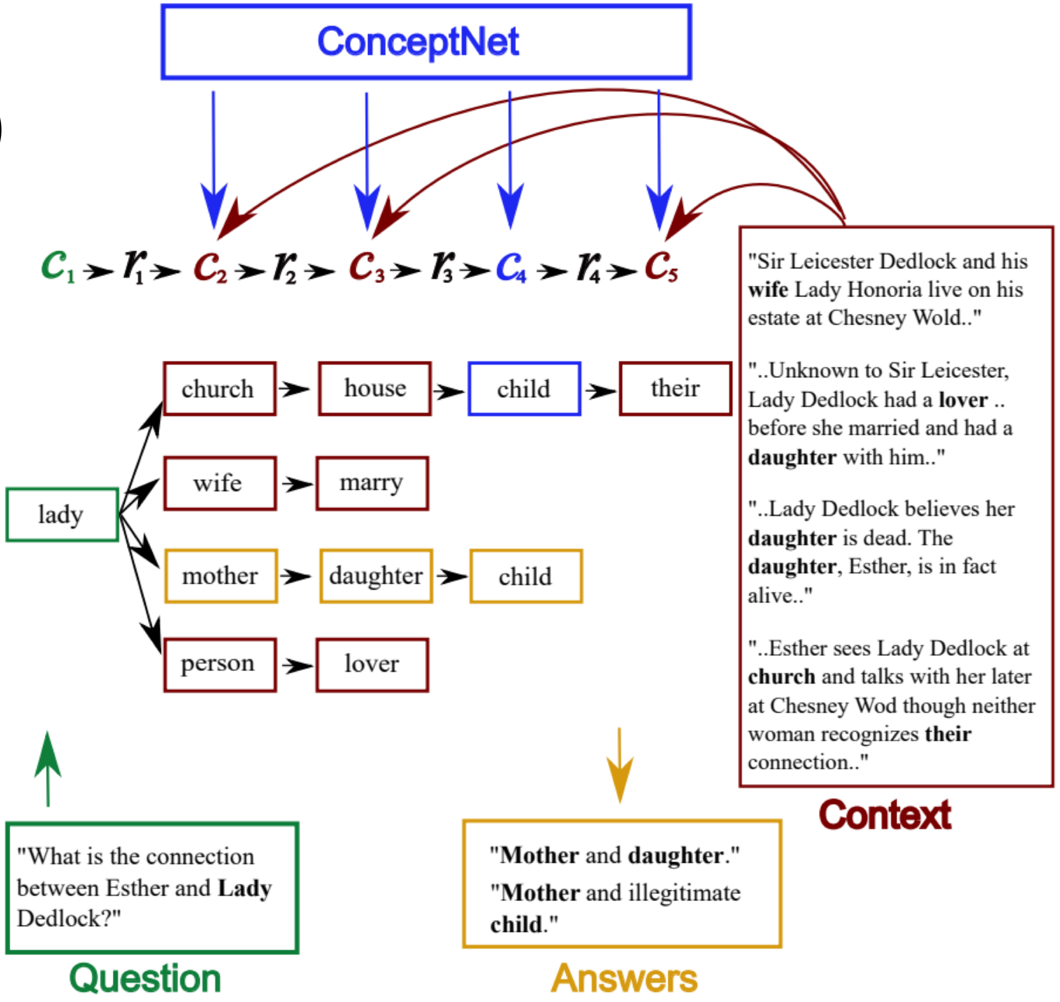

Bauer, L., Wang, Y., & Bansal, M. (2018). Commonsense for Generative Multi-Hop Question Answering Tasks. *EMNLP*.

Activation Value Visualisation for Question"What shore does Michael's corpse wash up on?"

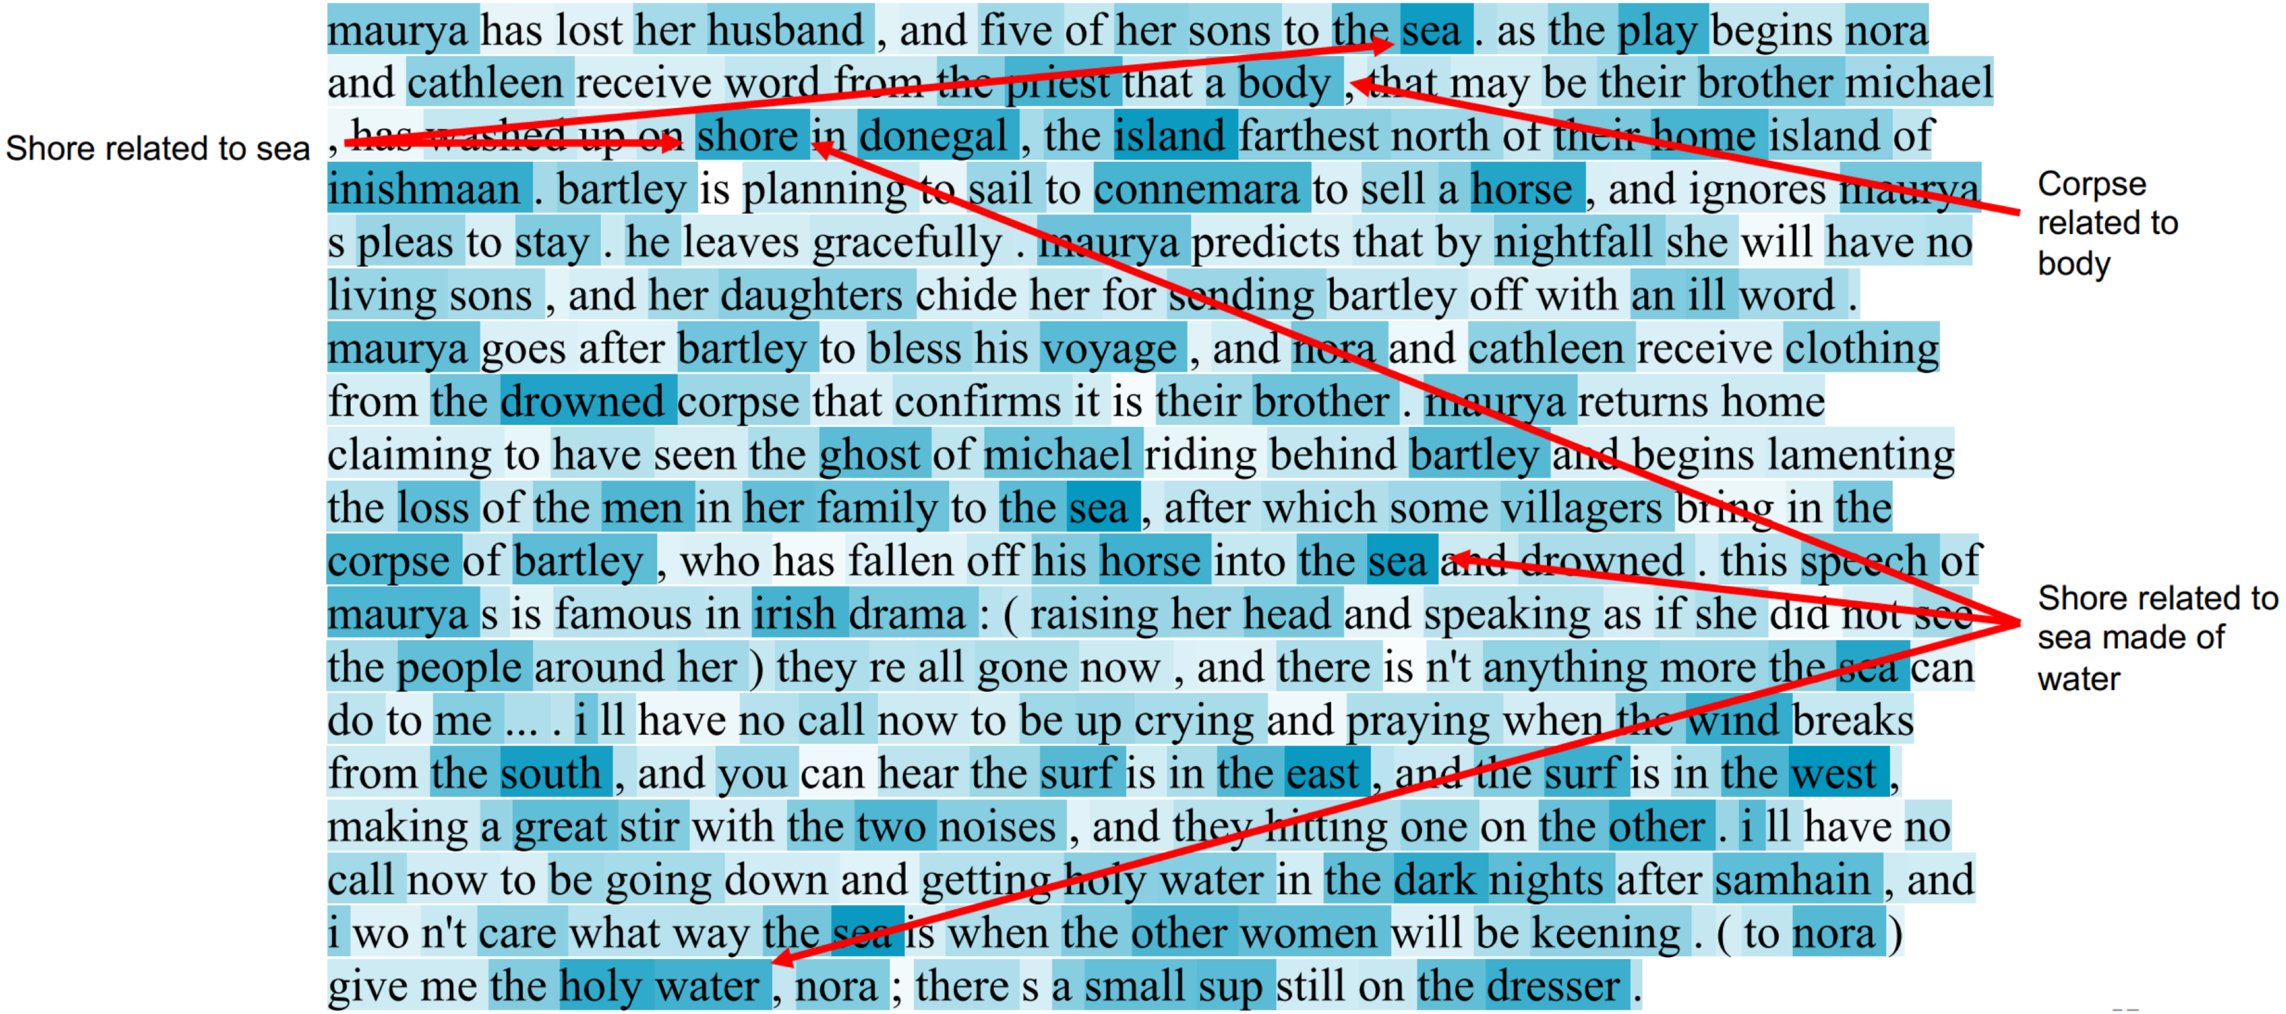

https://github.com/yicheng-w/CommonSenseMultiHopQA

Rationale Extraction

- Use continuous text span as the rationale to minimize the difficulty of the extraction task
- Introduce a rationale extraction task into the encoder
- Use a linear decay schedule to rely more on the rationale extraction task for addressing the semantic drift problem at the early stage

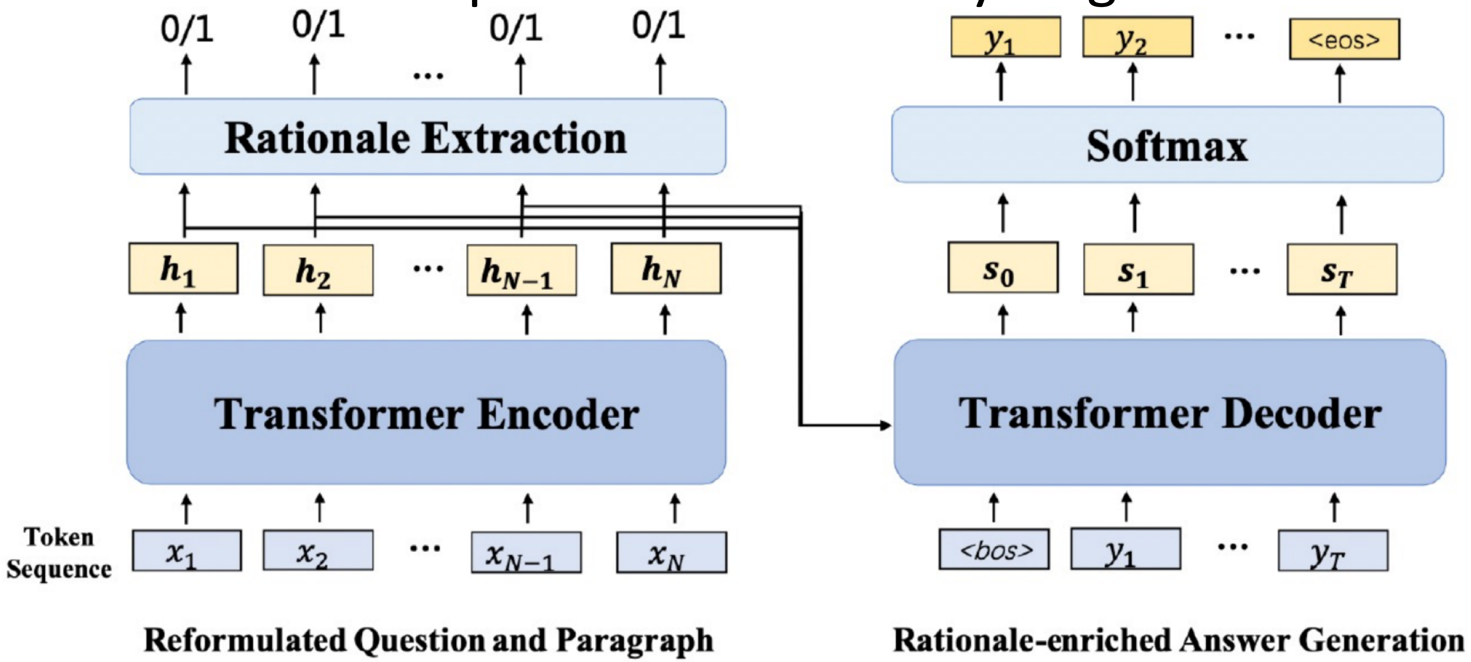

Li, C., Bi, B., Yan, M., Wang, W., & Huang, S. (2021). Addressing Semantic Drift in Generative Question Answering with Auxiliary Extraction. *ACL/IJCNLP*.

Example

Question	*Can a child get a flu vaccine under 6 months?*
Relevant Passage	*Yes | No Thank you!* ***Flu shots are not made for children under the age of 6 months.*** *If you read the vaccine insert and studies regarding the flu shot and kids, you will see that flu shots don't even work for children under the age of 2.*
Gold Answer	*No, a child under 6 months can't be given a flu vaccine.*
PALM Answer	*Yes, a child can get a flu vaccine under 6 months.*
REAG Answer	*No, a child cannot get a flu vaccine under 6 months.*

Li, C., Bi, B., Yan, M., Wang, W., & Huang, S. (2021). Addressing Semantic Drift in Generative Question Answering with Auxiliary Extraction. *ACL/IJCNLP*.

Application 4: Generative Commonsense Reasoning

- Problems with commonsense reasoning
  - Pre-trained language models are overly sensitive to co-occurrence
    - For example, consider a multi-choice question "*What do you fill with ink to write notes on a piece of copy paper? (A) fountain pen (B) pencil case (C) printer (D) notepad*", the pre-trained language model tends to predict *'(C) printer'*. The model may be overly sensitive to the co-occurrence between phrases in the question sentence like '*ink*' and '*copy paper*' and the answer choice 'printer
  - Generated sentences by pre-trained language models fail to capture commonsense
    - For example, given a set of commonsense concepts *"river, fish, net, catch"* , the GPT-2 generated "*A fish is catching in a net*"; UniLM generated "*A net catches fish*", etc.

Generative and Contrastive Objectives

External KG

Generative and Contrastive Objectives

Concept-to-Sentence

Input: <c2s> Generate a sentence with the concepts: forward, Simpson, ignore, information, prosecutor

Text-to-Text Transformer

Output: The information was forwarded to Simpson 's prosecutors, but it was ignored.

generative common sense

Concept Order Recovering

Input: <cor> Correct the order of the given sentence: Rahul stops him, fights his bar, and drives to a local performance.

Text-to-Text Transformer

Output:
Rahul fights him, stops his performance, and drives to a local bar.

Generative QA

Input: <cont> Which sentence is correct?: options:

1. The increased number of male visitors inspired by the article raised security concerns
2. The increased article of male visitors raised by the number inspired security concerns

discriminative common sense

Text-to-Text Transformer

Output:
The increased number of male visitors inspired by the article raised security concerns

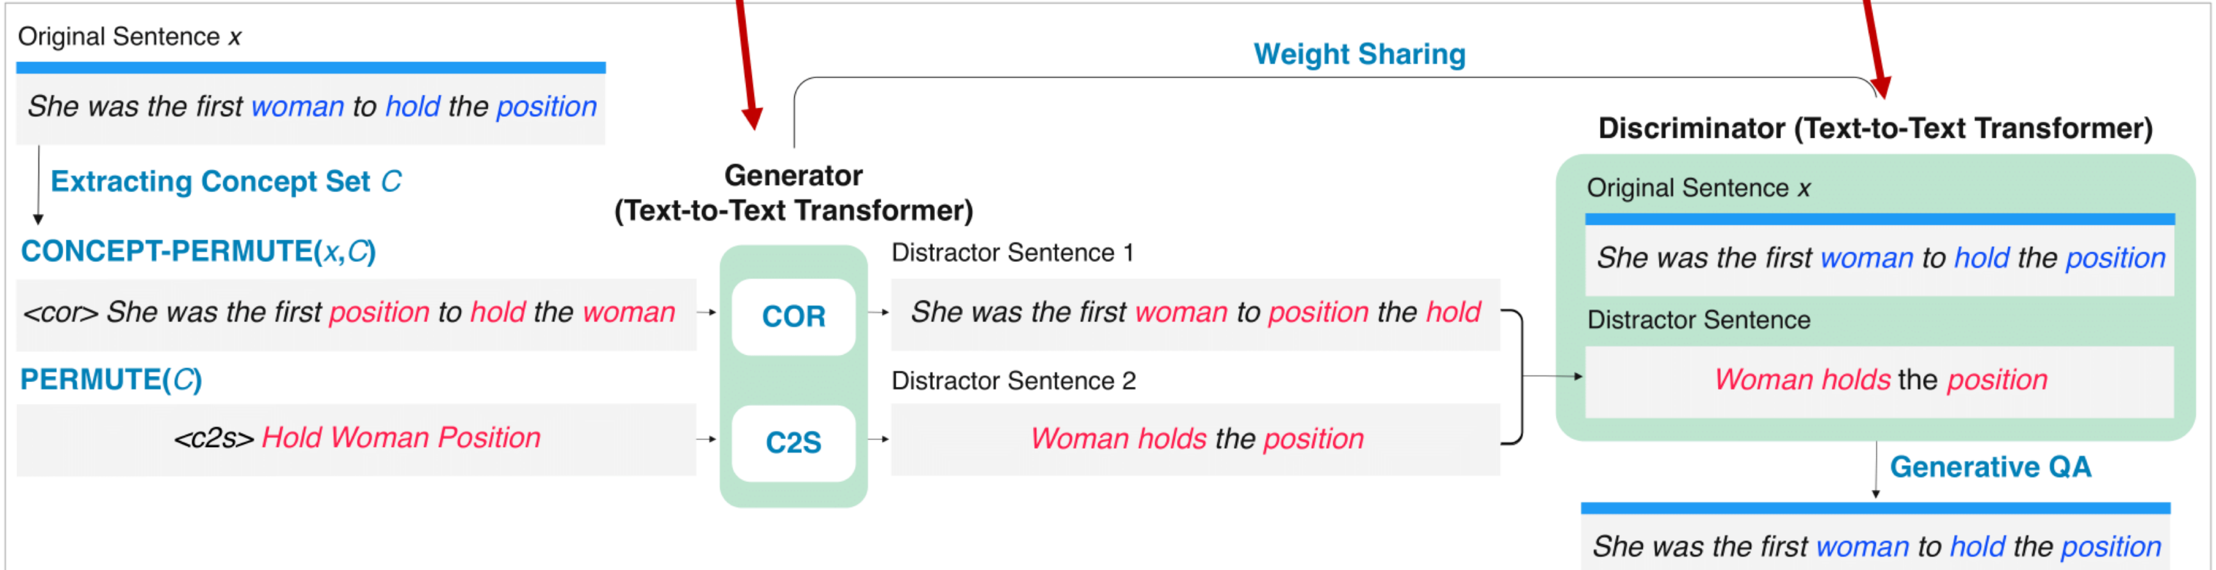

Zhou, W., Lee, D., Selvam, R.K., Lee, S., Lin, B., & Ren, X. (2021). Pre-training Text-to-Text Transformers for Concept-centric Common Sense. *ICLR*.

Example

Concept-Set	**T5-base**	**CALM-base**
Grass, Dog, Ball, Chase	*a dog is chased by a ball on the grass.*	*dog chasing a ball in the grass.*
Net, Cast, Boat, Water	*fishing boat casts a net in the water.*	*fisherman casts a net into the water from a fishing boat.*
Hole, Tree, Plant, Dig	*a man digs a hole in a tree to plant a new tree. he digs the*	*man digging a hole to plant a tree.*
Ingredient, Add, Pan, Fry	*a pan filled with ingredients adds a touch of spice to the fry.*	*add the ingredients to a pan and fry.*
Water, Hold, Hand, Walk	*A man holding a hand and walking in the water. A man is holding water.*	*man holding a bottle of water in his hand as he walks down the street.*
Place, Use, Metal tool	*A man uses a metal tool to make a piece of metal.*	*woman uses a metal tool to make a piece of jewelry.*

https://github.com/INK-USC/CALM

External KG

- Incorporate ConceptNet to both encoder and decoder
  - Ground concepts to ConceptNet
  - Use graph attention mechanism to integrate the entity-oriented knowledge information into token representations
  - Refine KG-augmented decoder with hierarchical graph attention mechanism

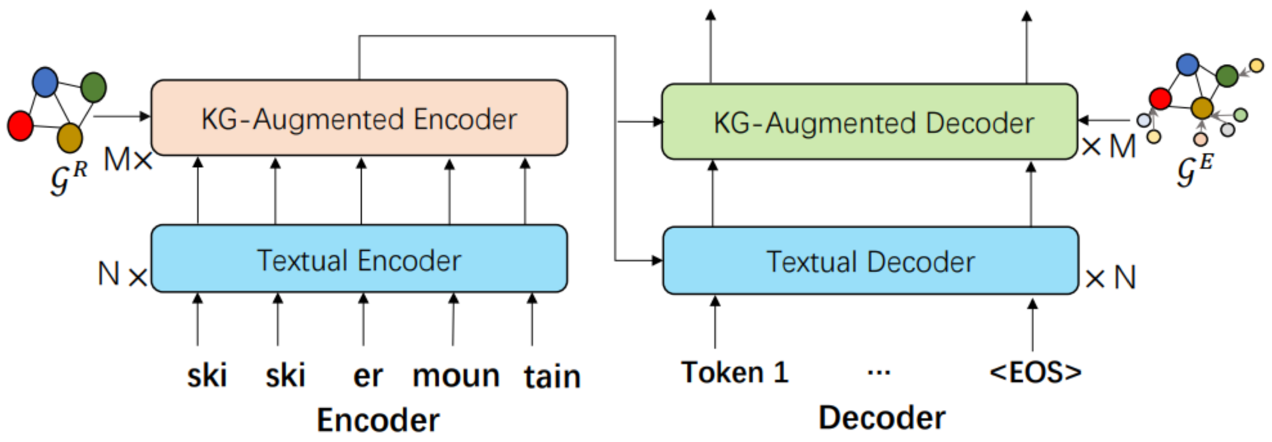

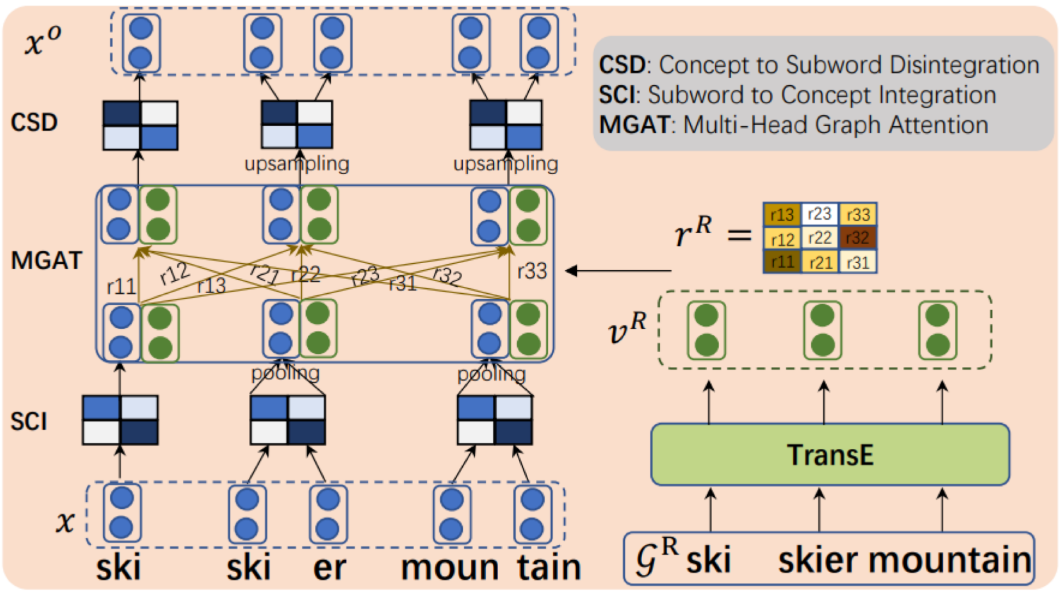

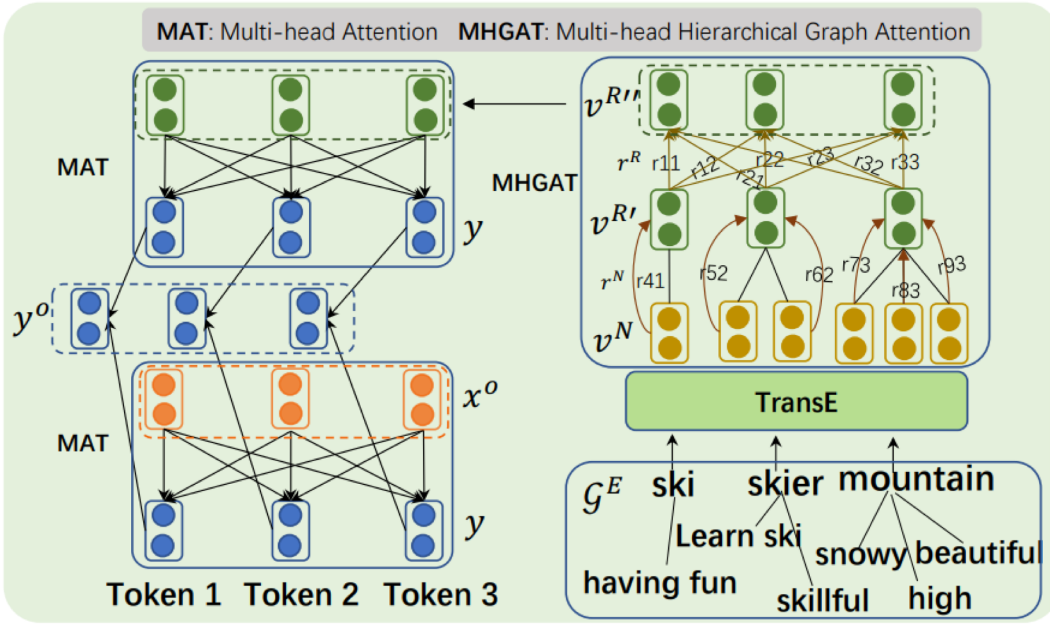

Liu, Y., Wan, Y., He, L., Peng, H., & Yu, P.S. (2021). KG-BART: Knowledge Graph-Augmented BART for Generative Commonsense Reasoning. *AAAI*.

Example

Concept Set: {stand, hold, street, umbrella }

[GPT-2]: A woman holding a umbrella in street
[BERT-Gen]: The woman stands on the street holding an umbrella.
[UniLM]: A man stands next to an umbrella on a street.
[T5]: A man holding an umbrella stands on a street.
[BART]: The woman holding an umbrella stands on the street and holds an umbrella.

1. A man held an umbrella while standing on the street.
2. People standing in the crowd street, many holding umbrellas.

[KG-BART]: A man holds an umbrella as he stands on the empty street.

https://github.com/yeliu918/KG-BART

Application 5: Story Generation

- Problems with story generation
  - Lack of knowledge
    - Generated story diverges from topics
    - Entity and event inconsistent with contexts
    - Lack of control granularity at sentence levels
  - Pre-trained language models suffer from inference capability
    - Cannot model causality

External KG and Keywords

Contextualized Inference Rules

External KG and Keywords

- Incorporate external knowledge base
- Allow controllability by replacing the keyword prediction process with manual external keywords

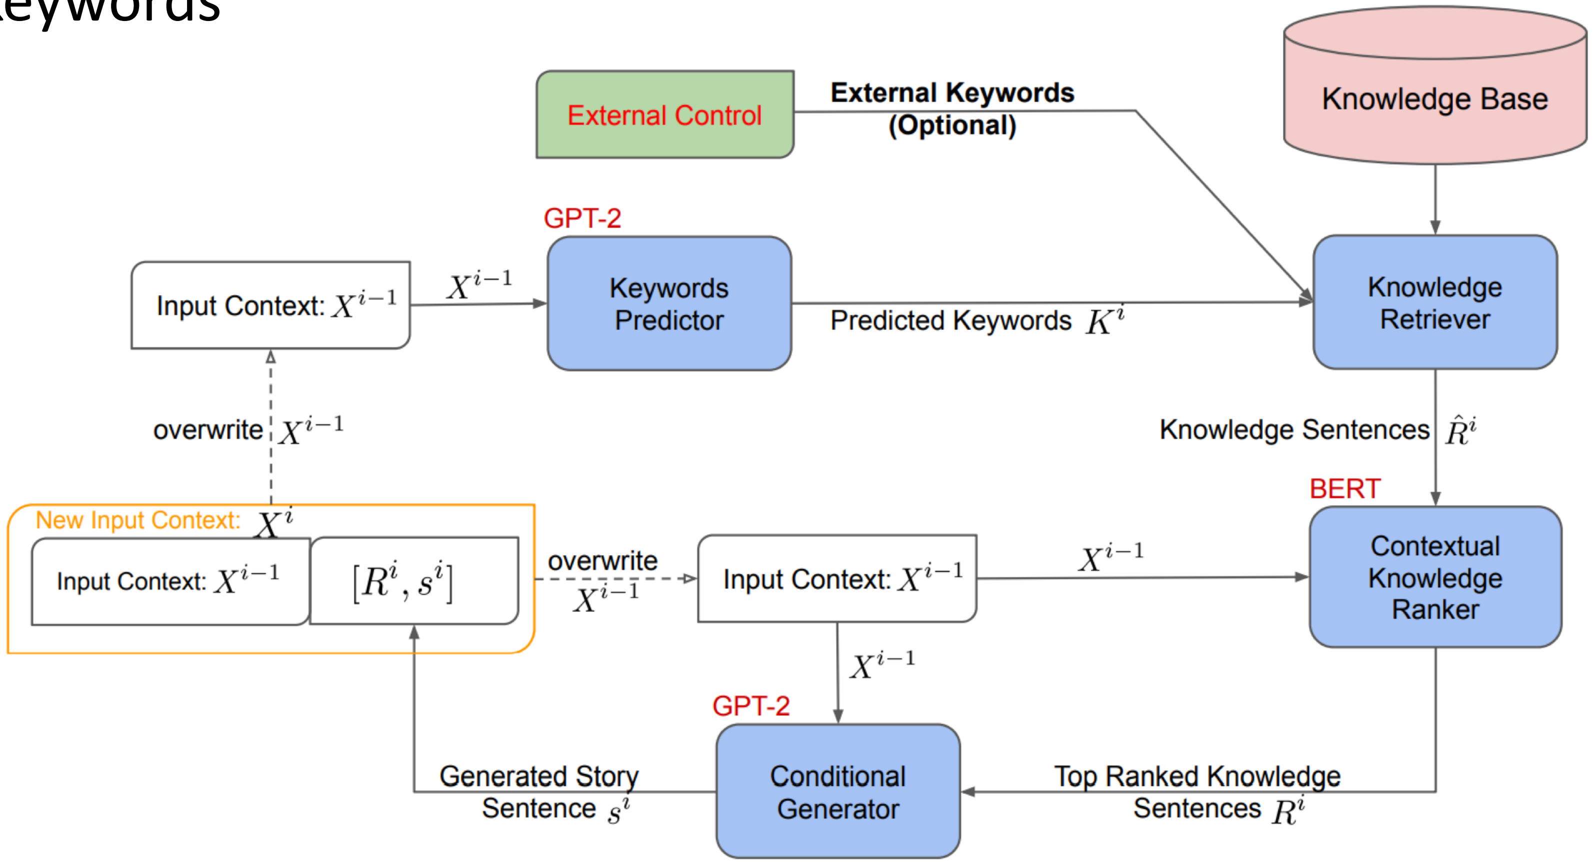

Xu, P., Patwary, M.A., Shoeybi, M., Puri, R., Fung, P., Anandkumar, A., & Catanzaro, B. (2020). Controllable Story Generation with External Knowledge Using Large-Scale Language Models. *EMNLP*.

Example

Context	*[MALE] and his friends are bored on a weekend day.*
Guan et al. (2019)	*they decided to sneak into a movie . they got a bunch of popcorn and a movie . then , they made it to a movie . now , they have a really good time .*
MEGATRON-CNTRL-124M	*keyword: go, decides* *it is going to be a big night. [MALE] and his friends decide to go out. they get drunk and have a great time. when they are done they feel happy.*
MEGATRON-CNTRL-8B	*keyword: go, decides* *they decided to go to the video game store. [MALE] picks out a new video game. [MALE] and his friends have fun playing the new video game. [MALE] is no longer bored.*
MEGATRON-CNTRL-8B-ANT	*keyword: come* *a new tv show begins to come on the local channel. [MALE] is excited to watch his favorite tv show. he watches his show and they all laugh together. [MALE] and his friends have a fun time watching their favorite show.*

https://github.com/NVIDIA/Megatron-LM

Contextualized Inference Rules

- Incorporate contextualized inference rules
  - Help model become more interpretable
  - Guide generation based on inference rules

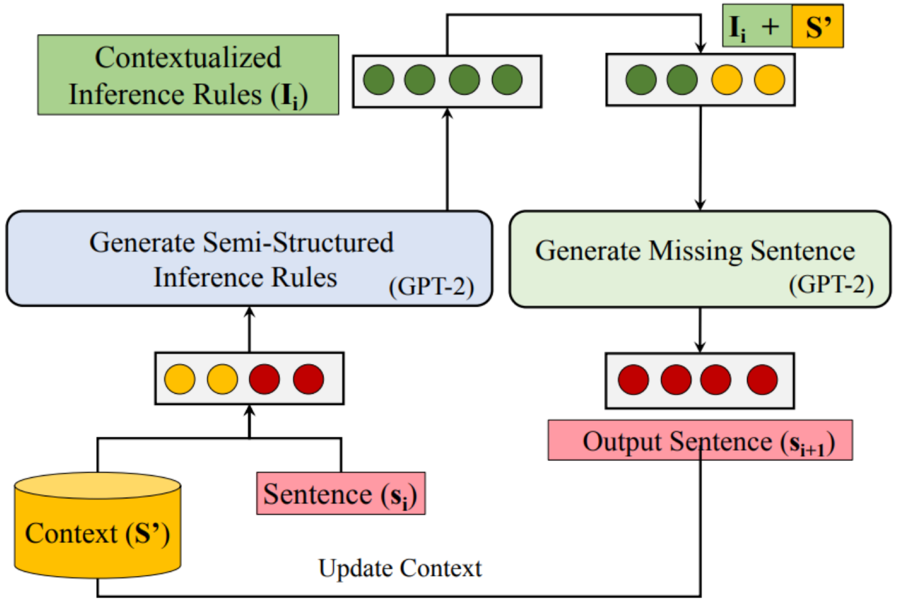

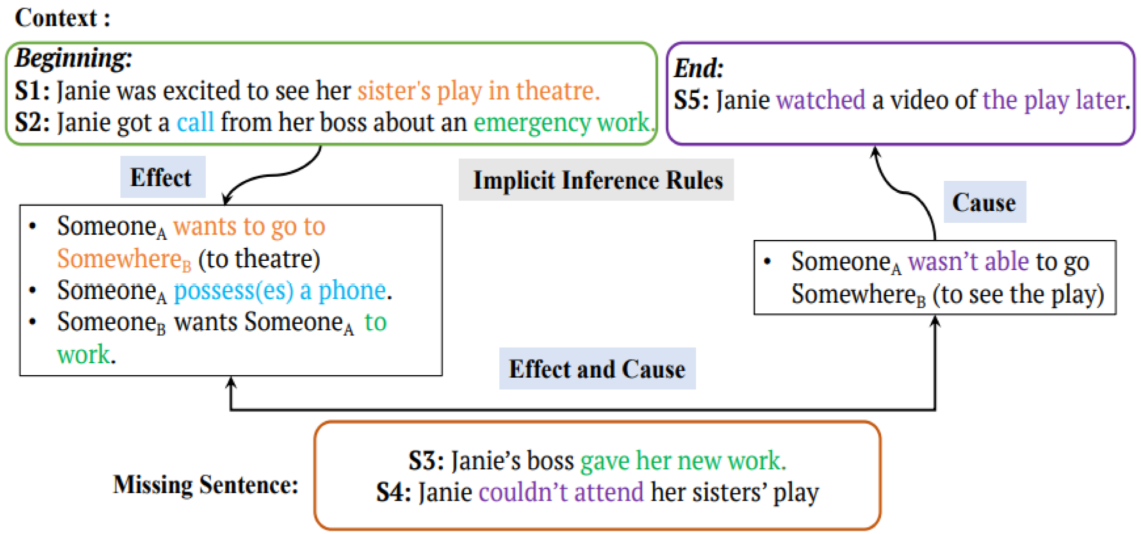

Paul, D., & Frank, A. (2021). COINS: Dynamically Generating COntextualized Inference Rules for Narrative Story Completion. *ACL*.

Example

Incomplete Story:	s_1: Ken was driving around in the snow. s_2: He needed to get home from work. s_5: His tires lost traction and he hit a tree.
Missing Sentences:	s_3: He was driving slowly to avoid accidents. s_4: Unfortunately the roads were too slick and Ken lost control.
COINS (I_{GR})	Someone$_A$ is going Somewhere$_B$ ≻Cause/Enables≻ Someone$_A$ is at Somewhere$_B$, Someone$_A$ **is driving Something$_A$ fast** ≻Cause/Enables≻ Something$_A$ hits Something$_B$ (that is a tree), Someone$_A$ possess(es) Something$_A$ (that is **a car**) ≻Enables≻> Something$_A$ (tires) lost Something$_B$ (traction)
COINS (I_{SR})	He posses(es) a car ≻result in≻ His tires lost traction, He needed to get home ≻Enables≻ He drove home, He was **driving on ice** ≻ Causes/Enables ≻ His tires lost traction, He was driving on ice ≻Causes/Enables≻ He **lost control** of his vehicle.
COINS(MS_{GR})	He was **driving too fast**. He lost control of his **car** .
COINS(MS_{SR})	He was **driving on ice**. He **lost control of his vehicle** .
GPT-2	He stopped at a gas station. He filled his tank.
GPT-2 GLUCOSE	When he got to the house he realized he was stuck. Ken had to pull over to get help.
KE	When he got home, he noticed his tires were flat. He decided to pull over.
GRF	He pulled over to see what was wrong. He saw that his car was stuck in the snow.
Human	He was going very fast. The street was slippery from the snow.

https://github.com/Heidelberg-NLP/COINS

Benchmarks

- Overview
- Evaluation
- Datasets
  - Dialog Systems
  - Question Answering
  - Question Generation
  - Commonsense Reasoning
  - Summarization

Benchmarks Overview

Tasks	Dataset Name	External Resources	Leaderboard	Benchmark
Dialog Systems	Wizard of Wikipedia	Wikipedia	Yes	Kilt
Question Answering	ELI5	Common Crawl	Yes	Kilt
Question Generation	SQuAD 1.1	Wikipedia	No	GLGE
Commonsense Reasoning	CommonGen	N/A	Yes	GEM
	αNLG-ART	N/A	Yes	GENIE
	ComVE	N/A	Yes	SemEval
Summarization	CNN/DailyMail	N/A	Yes	GLGE

Automatic Evaluation

- Untrained Automatic Metrics
  - N-Gram Overlap Metrics
    - ROUGE
    - METEOR
    - SacreBLEU
    - BLEU
    - GLGE-score
  - Fact-checking Metrics
    - PARENT
- Trained Automatic Metrics
  - Pre-trained Language Model-based Evaluation
    - BERTScore
    - BARTScore
    - BLEURT
    - QuestEval
    - KPQA

Human Evaluation

- Evaluation for Quality of Text
  - Fluency
  - Redundancy
  - Coherence
  - Commonsense
  - Grammar
  - Faithfulness
- Inter-Annotator Agreement
  - Percent agreement, Cohen’s Kappa, Fleiss’s Kappa, Krippendorff’s Alpha
- Problems
  - Expensive
  - Time-consuming
  - Quality Control
  - Challenge Criteria
  - Inconsistency in Evaluations
  - Inconsistency in Report

Dialog Systems : Wizard of Wikipedia

- An open-domain conversation dataset directly grounded with knowledge retrieved from Wikipedia
- Wizard
  - Inform their conversation partner (Apprentice) about a topic that one of them will choose
  - Have access to Wikipedia paragraphs that might be related to conversations
- Apprentice
  - Go into depth about a chosen topic that interests themselves or their partner (Wizard), while keeping the conversation engaging and fun

Dataset: https://parl.ai/projects/wizard_of_wikipedia/
Leaderboard: https://eval.ai/web/challenges/challenge-page/689/leaderboard/1909

Dialog Systems: Wizard of Wikipedia

Topic	*Lifeguard*
Apprentice	*So I am a lifeguard. Know anything about saving lives in water?*
Wizard	*I'm impressed! It's a big responsibility to supervise other people's safety in the water! Tell me more.*
Apprentice	*Well, I help make sure people do not drown or get injured while in or near the water!*
Knowledge	*A lifeguard is a rescuer who supervises the safety and rescue of swimmers, surfers, . . . Lifeguards are strong swimmers and trained in CPR/AED first aid, certified in water . . .* *. . .* *In some areas, the lifeguard service also carries out mountain rescues, or may function as the primary EMS provider*
Wizard	*I've heard that in some places, lifeguards also help with other sorts of emergencies, like mountain rescues!* *Is that part of your job too?*
Apprentice	*I have! I feel like you know much about this! What brings you to know so much?*
Wizard	*Oh, that's about the extent of my knowledge. I've just been around beaches and I've always admired lifeguards. I'm not a super strong swimmer myself.*

Question Answering: ELI5

- A long-form question answering dataset containing complex, diverse questions that require explanatory multi-sentence answers
- Use web search results (Common Crawl) as evidence documents to answer each question

Question	*How do Jellyfish function without brains or nervous systems?*
Supporting Documents	*The box jellyfish nervous system is divided into three functional parts namely; rhopalia, conducting nerve ring, and motor nerve net. [...] While they don't possess brains, the animals still have neurons that send all sorts of signals throughout their body. ``It is not true that jellyfish have no central nervous systems. They have an unusual nervous system,'' [...]*
Answer	*Jellyfish may not have a brain, but they have a rough nervous system and innate behaviours. However, they are very simple creatures. They're invertebrate: creatures without a backbone. Most jellyfish have really short life spans. [...] As their name implies, they are largely composed of basically jelly inside a thin membrane. They're over 95% water.*

Dataset and Leaderboard: https://facebookresearch.github.io/ELI5/

Question Generation: SQuAD 1.1

- A reading comprehension dataset consisting of 100,000+ questions posed by crowd workers on a set of Wikipedia articles
- The answer to each question is a segment of text from the corresponding reading passage

Passage	*Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season. The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers 24–10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California. As this was the 50th Super Bowl, [...].*
Answer	*Santa Clara, California*
Target	*Where did Super Bowl 50 take place?*

Dataset: https://rajpurkar.github.io/SQuAD-explorer/

Commonsense Reasoning: CommonGen

- A constrained text generation task, associated with a benchmark dataset
- Explicitly test machines for the ability of generative commonsense reasoning
- Based on visually-grounded sentences from several existing caption datasets

Dataset and Leaderboard: https://inklab.usc.edu/CommonGen/

Commonsense Reasoning: αNLG-ART

- A generative commonsense reasoning dataset consists of over 20k commonsense narrative contexts and 200k explanations
- Given the observations at time t1 and t2, the model needs to generate a plausible hypothesis

Observation at t1	*Larry's yard was covered in dead leaves.*
Observation at t2 (t2>t1)	*Larry decided to give up for the day and went back inside.*
Hypothesis	*He spent hours trying to clean the yard.*

Dataset and Leaderboard: http://abductivecommonsense.xyz/

Commonsense Reasoning: ComVE

- Generate the reason why a statement is against common sense and use BELU to evaluate it
- Consists 2021 against common-sense sentences with true reasons

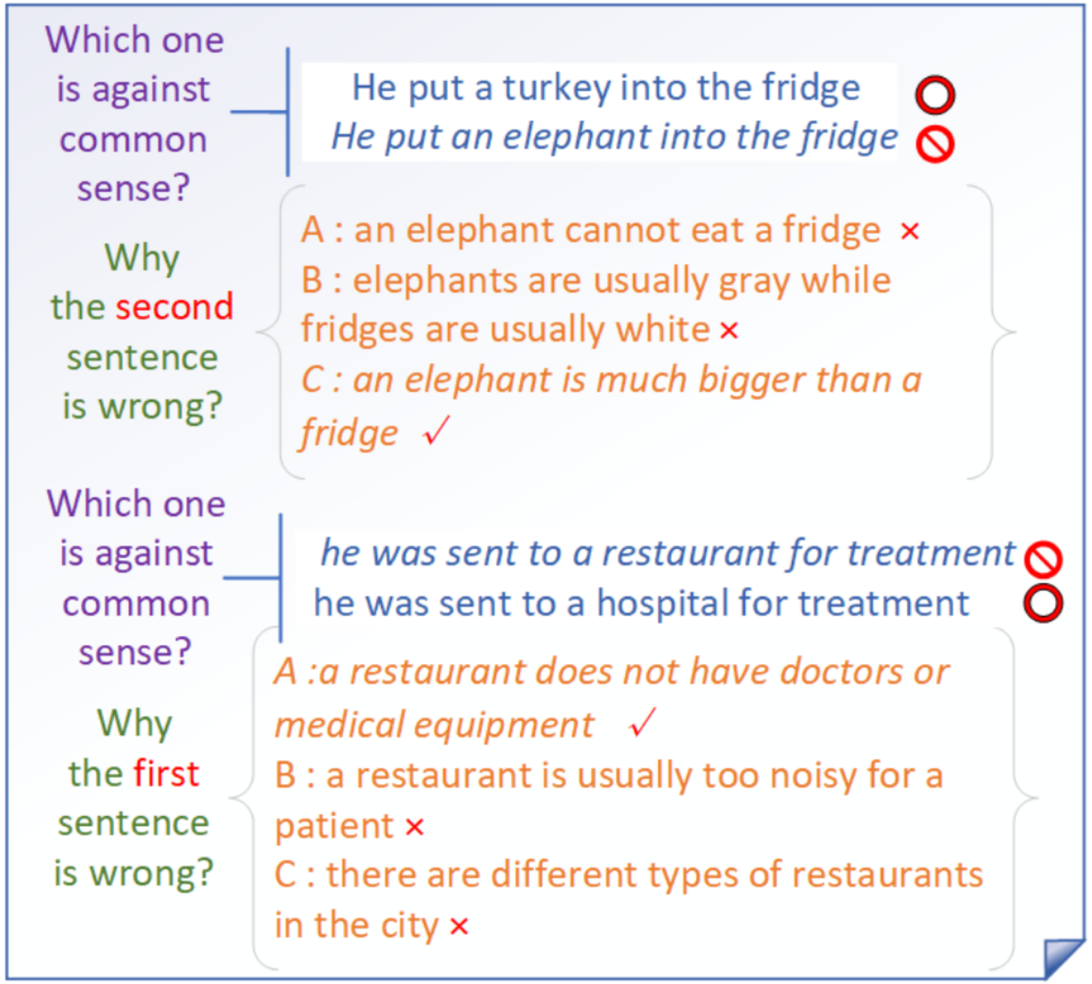

Dataset and Leaderboard: https://competitions.codalab.org/competitions/21080

Summarization: CNN/DailyMail

- A non-anonymized variant of CNN/DailyMail dataset consists of 311,971 <article, summary> pairs

Article	*andy murray came close to giving himself some extra preparation time for his wedding next week before ensuring that he still has unfinished tennis business to attend to. the world no 4 is into the semi-finals of the miami open, but not before getting a scare from 21 year-old austrian dominic thiem, who pushed him to 4-4 in the second set before going down 3-6 6-4, 6-1 in an hour and three quarters. murray was awaiting the winner from the last eight match between tomas berdych and argentina's juan monaco. prior to this tournament thiem lost in the second round of a challenger event to soon-to-be new brit aljaz bedene. andy murray pumps his first after defeating dominic thiem to reach the miami open semi finals. muray throws his sweatband into the crowd after completing a 3-6, 6-4, 6-1 victory in florida. murray shakes hands with thiem who he described as a 'strong guy' after the game. (…)*
Summary	*british no 1 defeated dominic thiem in miami open quarter finals. andy murray celebrated his 500th career win in the previous round. third seed will play the winner of tomas berdych and juan monaco in the semi finals of the atp masters 1000 event in key biscayne*

Dataset : https://github.com/becxer/cnn-dailymail/

Hands-on for GeDi Generative Discriminator Guided Sequence Generation

- We will test the model from GeDi: Generative Discriminator Guided Sequence Generation using Google Colab.
- GeDi guides generation at each step by computing classification probabilities for all possible next tokens via Bayes rule by normalizing over two class-conditional distributions
  - one conditioned on the desired attribute, or control code,
  - another conditioned on the undesired attribute, or anti control code

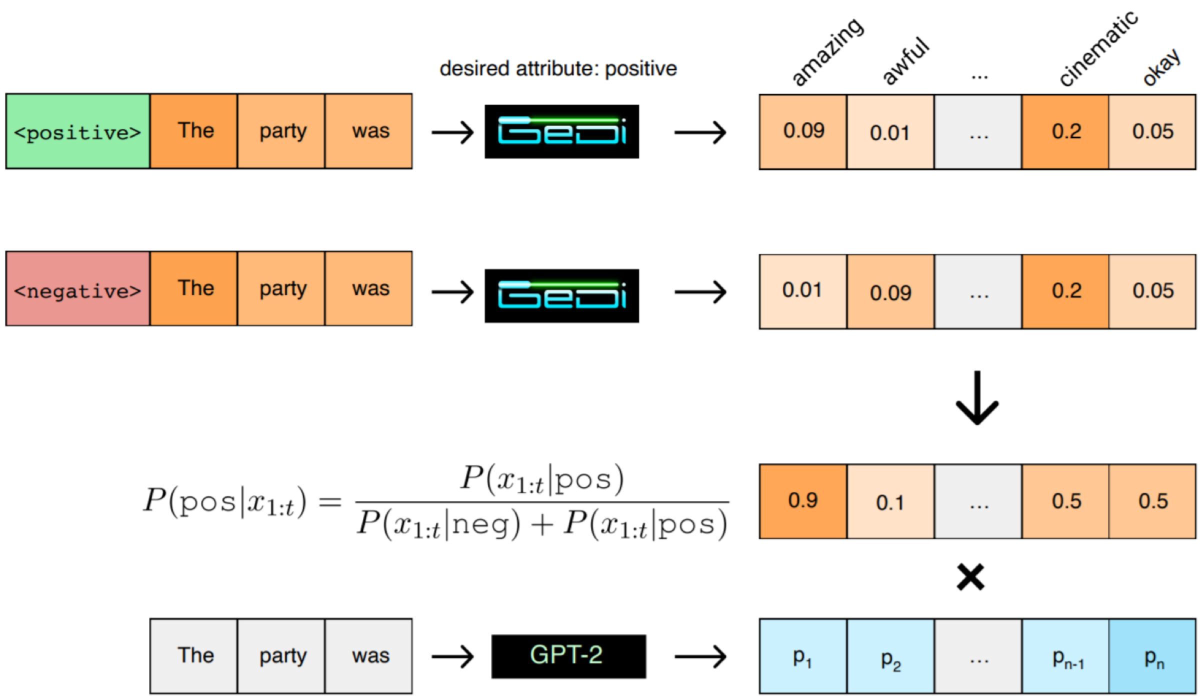

Krause, B., Gotmare, A., McCann, B., Keskar, N.S., Joty, S.R., Socher, R., & Rajani, N. (2020). GeDi: Generative Discriminator Guided Sequence Generation. *ArXiv, abs/2009.06367*.

Download Code Repository

- Companion Notebook by Salesforce is here: https://colab.research.google.com/github/salesforce/GeDi/blob/master/GeDi_guided_GPT_2_XL.ipynb

```
[1]  !wget https://storage.googleapis.com/sfr-gedi-data/GeDi.zip
     import zipfile
     with zipfile.ZipFile('GeDi.zip', 'r') as zip_ref:
         zip_ref.extractall('./')
```

```
--2021-11-05 23:58:09--  https://storage.googleapis.com/sfr-gedi-data/GeDi.zip
Resolving storage.googleapis.com (storage.googleapis.com)... 64.233.191.128, 209.85.145.128, 172.217.219.128, ...
Connecting to storage.googleapis.com (storage.googleapis.com)|64.233.191.128|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 260070 (254K) [application/zip]
Saving to: 'GeDi.zip'

GeDi.zip            100%[===================>] 253.97K  --.-KB/s    in 0.003s

2021-11-05 23:58:09 (92.8 MB/s) - 'GeDi.zip' saved [260070/260070]
```

Install Dependencies

```
[3]  '''Installing transformers v2.8'''

     !pip install transformers==2.8
     !pip install -r hf_requirements.txt

     '''Downloading GeDi topic model checkpoints'''

     !wget https://storage.googleapis.com/sfr-gedi-data/gedi_topic.zip

     with zipfile.ZipFile('gedi_topic.zip', 'r') as zip_ref:
         zip_ref.extractall('./')
```

```
Collecting transformers==2.8
  Downloading transformers-2.8.0-py3-none-any.whl (563 kB)
     |████████████████████████████████| 563 kB 5.5 MB/s
Collecting sacremoses
  Downloading sacremoses-0.0.46-py3-none-any.whl (895 kB)
     |████████████████████████████████| 895 kB 38.0 MB/s
Requirement already satisfied: requests in /usr/local/lib/python3.7/dist-packages (from transformers==2.8) (2.23.0)
Collecting tokenizers==0.5.2
  Downloading tokenizers-0.5.2-cp37-cp37m-manylinux1_x86_64.whl (5.6 MB)
     |████████████████████████████████| 5.6 MB 26.4 MB/s
Requirement already satisfied: regex!=2019.12.17 in /usr/local/lib/python3.7/dist-packages (from transformers==2.8) (2019.12.20)
Requirement already satisfied: filelock in /usr/local/lib/python3.7/dist-packages (from transformers==2.8) (3.3.0)
Collecting boto3
```

Download GPT2-XL Tokenizer

```
[4] import numpy as np
    import torch


    from modeling_gpt2 import GPT2LMHeadModel

    from transformers import (
        GPT2Config,
        GPT2Tokenizer
    )
```

```
[5] mode = "topic"
    code_desired = "true"
    code_undesired = "false"
    model_type = 'gpt2'
    gen_type = "gedi"
    gen_model_name_or_path = "gpt2-xl"

    device = torch.device("cuda" if torch.cuda.is_available() else "cpu")

    MODEL_CLASSES = {"gpt2": (GPT2Config, GPT2LMHeadModel, GPT2Tokenizer),}
    config_class, model_class, tokenizer_class = MODEL_CLASSES["gpt2"]
    tokenizer = tokenizer_class.from_pretrained(gen_model_name_or_path, do_lower_case=False)
```

```
Downloading: 100% 1.04M/1.04M [00:00<00:00, 1.00MB/s]
Downloading: 100% 456k/456k [00:00<00:00, 993kB/s]
```

Download and Covert GPT2-XL

```
[6] #Loading GPT2-XL model (1.5B param LM) below, this could take a while.
    #This requires additional CPU memory overhead to load the pretrained weights in a new model
    #Due to CPU memory constraints on Colab, we're loading the model in half precision (load_in_half_prec=True)
    #Do to this change, generations may not always exactly match samples in paper, but sometimes do, and seem to be similar in quality
    #If you run the notebook with enough CPU RAM (most likely 16GB+), you can try setting load_in_half_prec=False

    model = model_class.from_pretrained(gen_model_name_or_path, load_in_half_prec=True)
    model = model.to(device)
    model = model.float()

    gedi_model_name_or_path = 'gedi_topic'
    gedi_model = model_class.from_pretrained(gedi_model_name_or_path)
    gedi_model.to(device)
```

```
Downloading: 100% 787/787 [00:00<00:00, 10.2kB/s]
Downloading: 100% 6.43G/6.43G [03:48<00:00, 30.2MB/s]
```

Set Arguments for Generation

```
[7] #setting arguments for generation
    #max generation length
    gen_length = 200
    #omega from paper, higher disc_weight means more aggressive topic steering
    disc_weight = 30
    #1 - rho from paper, should be between 0 and 1 higher filter_p means more aggressive topic steering
    filter_p = 0.8
    #tau from paper, preserves tokens that are classified as correct topic
    target_p = 0.8
    #hyperparameter that determines class prior, set to uniform by default
    class_bias = 0

    if gen_length>1024:
      length = 1024
    else:
      length = gen_length
```

Specify Prompt and Topic to GeDi

```
[8] #Specify what topic you want to generate on using the secondary_code variable

    secondary_code = 'climate'
    bpe_tokens = tokenizer.encode(secondary_code)
    if len(bpe_tokens) > 1:
      print("Warning! number of bpe tokens for " + code + " is greater than 1, model isn't trained for this, generation is less likely to match the topic")
```

```
[9] #Specify prompt below
    prompt = "In a shocking finding"

    start_len=0
    text_ids = tokenizer.encode(prompt)
    encoded_prompts=torch.LongTensor(text_ids).unsqueeze(0).to(device)

    multi_code = tokenizer.encode(secondary_code)
    attr_class = 1
```

Predict Results

```
generated_sequence = model.generate(input_ids=encoded_prompts,
                                         pad_lens=None,
                                          max_length= length,
                                          top_k=None,
                                          top_p=None,
                                          repetition_penalty= 1.2,
                                          rep_penalty_scale= 10,
                                          eos_token_ids = tokenizer.eos_token_id,
                                          pad_token_id = 0,
                                          do_sample= False,
                                          penalize_cond= True,
                                          gedi_model= gedi_model,
                                          tokenizer= tokenizer,
                                          disc_weight= disc_weight,
                                          filter_p = filter_p,
                                          target_p = target_p,
                                          class_bias = class_bias,
                                          attr_class = attr_class,
                                          code_0 = code_undesired,
                                          code_1 = code_desired,
                                          multi_code=multi_code
                                          )

text = tokenizer.decode(generated_sequence.tolist()[0], clean_up_tokenization_spaces=True)
print('\n')
print(text)
```

```
GeDi estimates the probability that it sample is desired class is: 0.5867899656295776


In a shocking finding that scientists are calling 'extremely worrying', the world's oceans are becoming increasingly acidic.

According to new research published in Nature Climate Change, ocean waters around the world are becoming significantly more acidic due to rising level

"Our results show that ocean acidification has already begun in many regions, with most regions experiencing acidification rates greater than predicte


Ocean Acidification is an environmental change caused by increases in atmospheric carbon dioxide (CO2), resulting in increased concentrations of CO2<|
```

Output Results

- GeDi estimates the probability that it sample is desired class is: 0.5867899656295776
- *In a shocking finding that scientists are calling 'extremely worrying', the world's oceans are becoming increasingly acidic.*
- *According to new research published in Nature Climate Change, ocean waters around the world are becoming significantly more acidic due to rising levels of carbon dioxide (CO2) in the atmosphere.*
- *"Our results show that ocean acidification has already begun in many regions, with most regions experiencing acidification rates greater than predicted for preindustrial conditions by 2100," says lead author Thomas Crowley, who conducted this research as part of his doctoral degree thesis at The University of Western Australia. "Ocean acidification has important consequences for organisms living near or below sea surface because low pH environments may be particularly challenging for calcifying organisms; however, our results also show that marine ecosystems will likely experience increasing acidification rates even when they don't experience current ocean acidity trends."*
- *Ocean Acidification is an environmental change caused by increases in atmospheric carbon dioxide (CO2), resulting in increased concentrations of CO2<|endoftext|>*

Another Related Project on Colab

- Tune GPT2 to generate controlled sentiment reviews
- Companion Notebook by Leandro von Werra is here: https://colab.research.google.com/github/lvwerra/trl/blob/master/nbs/05-gpt2-sentiment-control.ipynb
- The project optimizes GPT2 to produce IMDB movie reviews with controlled sentiments using a BERT sentiment classifier for rewards
- The project website is here