



CTRL

A CONDITIONAL TRANSFORMER LANGUAGE MODEL FOR CONTROLLABLE GENERATION

NITISH SHIRISH KESKAR , BRYAN MCCANN, LAV R. VARSHNEY, CAIMING XIONG, RICHARD SOCHER

Presented By - Karthik Ramesh

CS886

University of Waterloo

March 9, 2020

MOTIVATION

- Language models are often trained as conditional language models on specific tasks as required
- Even an unconditional language model is transferred to required tasks through transfer learning
- Unconstrained generation of text is not very well understood

“Inspired by the degree of control available in image generation..... we train a language model ... that make desired features of generated text more explicit”

OVERVIEW

- They aim to preserve the generality of the model through the usage of specific “control codes”
- These control codes are pre-defined and correspond to a specific domain (eg: Wikipedia, Amazon Reviews)
- The intuition is that using these control codes can allow the language model to trace back to the language structure that it had seen during training and base it’s generation on those sentences during inference time

MODEL

- Basically uses the decoder from the transformer architecture to construct the language model
- It is trained from scratch using the data along with conditional codes

“It uses 48 layers and 16 heads. Model has 1.6B parameters.”

TRAINING PHASE

- Control Codes are embedded into the training phrases depending on the type of control code
- Domain control codes are propagated to all text in the domain as the first token

Wikipedia *Anarchism* is a political philosophy that advocates the abolition of all forms of hierarchy
Books *Anarchism* is the \n only true and practical form of Socialism. It has been said that Socialism

- Other control codes are appended as necessary

Questions *Q: What is the capital of India?* *A:* New Delhi\n\n*Q: Which city in India has a large*
English : I lost 10 kgs! ; *German :* Ich habe 10 Kilogramm verloren!\n\n*English :* It is

GENERATION PHASE - SAMPLING

Sampling:

- In the equation above, $T \rightarrow 0$ approximates a greedy distribution which magnifies the peaks in the probability distribution while $T \rightarrow \infty$ flattens the distribution to make it more uniform

$$p_i = \frac{\exp(x_i/T)}{\sum_j \exp(x_j/T)}$$

- They modified the temperature-controlled sampling to get the below stated equation

- They proposed to use $\theta \approx 1.2$ and greedy sampling
- This θ is a penalty for repeated sampling of same token

$$p_i = \frac{\exp(x_i/(T \cdot I(i \in g)))}{\sum_j \exp(x_j/(T \cdot I(j \in g)))}$$

$$I(c) = \theta \text{ if } c \text{ is True else } 1$$

GENERATION PHASE – CONTROL CODES

CONTROL CODES:

- Style by domain – constraining to a particular domain
- More complex control codes – further constraining the generation phase; eg: a specific rating value
- Triggering specific tasks – triggering the Q&A in the model and translation
- Zero-shot code-mixing – mixing codes not seen before in the training data

RESULTS – STYLE BY DOMAIN

Horror *A knife* handle pulled through the open hole in the front. I jumped when the knife hit.\n\nEyes widened in horror. Her scream was the only sound I heard besides her sobs.\n\nThe spider touched her feet as it started to dig into the top of her arch. The creature's tears began to flow.\n\nThe spider looked up to her and looked back at me with eyes filled with burning tears. My heart started to race. . .

- For the same conditional text (prompt), the model has variation in generated text

Reviews *A knife* is a tool and this one does the job well.\n\nRating: 4.0\n\nI bought these for my husband who has been using them to cut up his own meat since he got them. He says they are very sharp so be careful when you use them, but that doesn't seem like much of an issue because he's used it on everything from chicken breasts to beef tenderloin. . .

RESULTS – COMPLEX CONTROL CODES

Reviews Rating: 5.0\n\n I have been using this product for a few years and it is the best thing on my hair. It makes your hair soft but not greasy.\n\nRating: 5.0\n\nI've used this product before and love it. It's light weight yet keeps my curls defined without making them crunchy!

- The generated text has been further constrained due to a specific rating
- Generates valid text without any type of prompt

Reviews Rating: 1.0\n\nI bought this for my daughter and it was a complete waste of money. The case is very flimsy and the screen protector doesn't fit properly.\n\nRating: 1.0\n\nThis product did not work at all. It would charge but when you unplug it, it does nothing. Very disappointed.

RESULTS – TRIGGERING SPECIFIC TASKS

Questions *Q: What is the capital of India?* **A:** New Delhi
Q: Which city in India has a large population of people from Gujarat?
A: Ahmedabad
Q: How many languages are spoken by Indians living abroad?
A: more than 100. . .

- Able to perform translation task without being explicitly trained to do so

Translation **English :** *We release a new model for coherent language generation ;* **French :** *Nous publions un nouveau modèle de génération cohérente du langage*
English : The European Parliament has just adopted the report on the proposal for a regulation of the European Parliament and of the Council establishing an action programme to promote multilingualism in Europe. ; **French :** Le Parlement européen vient d'adopter le rapport sur la proposition de règlement du Parlement et du Conseil établissant un programme-cadre pour promouvoir le plurilinguisme dans l'Union européenne. . .

RESULTS – ZERO SHOT CODE MIXING

- Able to combine the Diet domain along with translation task

Diet English : *I lost 10 kgs!* ; German : Ich habe 10 Kilogramm verloren!\n\nEnglish : It is a great achievement for me and my family. ; German : Das ist eine groe Leistung fr mich und meine Familie.\n\nEnglish : Thank you very much to all the people who have helped me. ; German : Vielen Dank an alle, die mir geholfen haben...

- Politics domain is mixed with French prompts despite this never occurring in the training data

Politics Title: *Les Etats-Unis sont un* pays de droite\n\nText: Je suis Américain et je ne peux pas supporter ce que font les USA en matière d’immigration\n\nJ’ai déjà écrit sur le sujet mais il n’y a aucune réaction http://www.reddit.com/r/politics/comments/1tqjzw/les_etats_unes_son_tres-de-gauche/ ...

RESULTS

QUANTITATIVE EVALUATION?

Moving onto PPLM

PPLM

PLUG AND PLAY LANGUAGE MODELS: A SIMPLE APPROACH TO CONTROLLED TEXT GENERATION

SUMANTH DATHATHRI, ANDREA MADOTTO, JANICE LAN, JANE HUNG, ERIC FRANK, PIERO MOLINO, JASON YOSINSKI, ROSANNE LIU

MOTIVATION

- Guide the generation of the language model with minimum modification in the model itself
- This enable us to maintain the generality of the language model while allowing us to generate the text with the desired attribute
- This is a more hands-off approach compared to CTRL

OVERVIEW

- This paper proposes to generate text with the required attribute using the language model by just appending an “attribute classifier” and using this to guide the generation of text
- So, this first involves building an attribute classifier whose input feature is the “condensed output” of the language model and training it to identify the required attribute based on this
- So now, to push the LM’s generated text towards the direction of the required attribute, we just backprop the gradients from the attribute classifier to modify the (K,V) pairs of every input till the current time-step

Uses a pre-trained GPT-2 model as the LM

TRAINING PHASE


ASCENDING $\log p(a|x)$

- Move the (K,V) in the direction that increases the likelihood of that attribute
- This provides the control over the generated text

$$\Delta H_t \leftarrow \Delta H_t + \alpha \frac{\nabla_{\Delta H_t} \log p(a|H_t + \Delta H_t)}{\|\nabla_{\Delta H_t} \log p(a|H_t + \Delta H_t)\|^\gamma}$$

ASCENDING $\log p(x)$

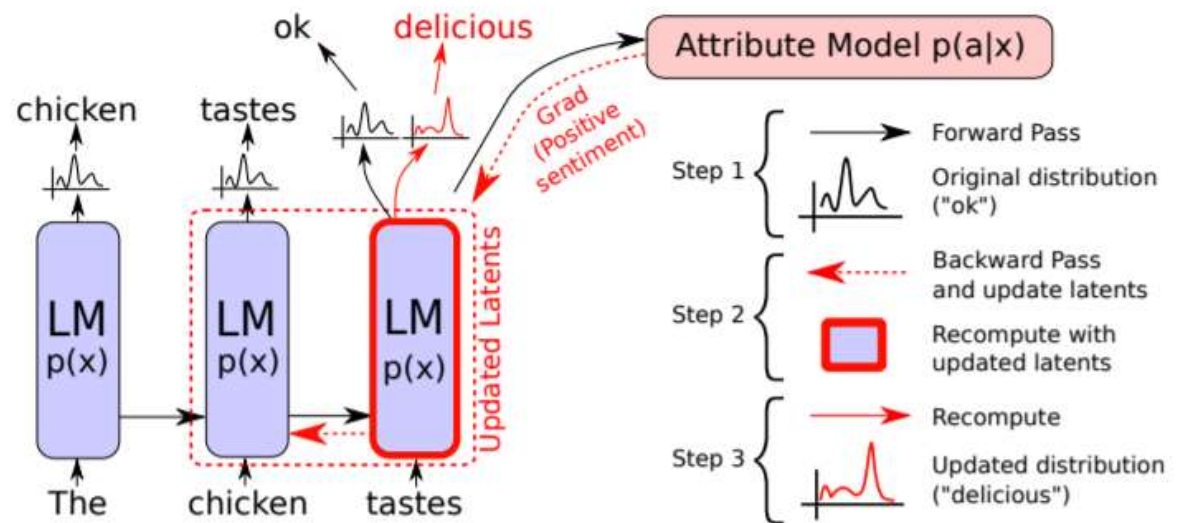
- Move the (K,V) pairs in the direction of the original distribution
- This ensures that the model does not end up generating text from the low probability regions and to maintain fluency



TRAINING PHASE

Model follows 3 steps: -

- Forward Pass through the model
- Backward pass to update the (K, V)
- Recomputing the output of the model using the updated latents



RESULTS – EVALUATION CRITERIA

- Automated Evaluation:
 - Perplexity is measured using GPT
 - Dist-1, Dist-2, Dist-3 scores -> they measure the number of distinct 1,2,3-grams across all the samples
- Human Evaluation:
 - Fluency – rank on a scale of 1-5
 - A/B Testing – rank the given pair on the topic relevance, sentiment strength

RESULTS – EVALUATION CRITERIA

They consider 4 scenarios for evaluation:

- **B** – The unchanged GPT-2 model sampled once
- **BR** – **B** but with sampling r times
- **BC** – updating the latent representations
- **BCR** – **BC** but sampled r times

All of the samples are filtered based on Dist scores and the best sample according to the LL is chosen

RESULTS

- BC and BCR do give a good amount of control compared to the rest
- Perplexity and fluency does degrade in BC and BCR but the Dist scores are maintained

Method	Topic % (↑ better) (human)	Perplexity (↓ better)	Dist-1 (↑ better)	Dist-2 (↑ better)	Dist-3 (↑ better)	Fluency (↑ better) (human)
B	11.1	39.85±35.9	0.37	0.79	0.93	3.60±0.82
BR	15.8	38.39±27.14	0.38	0.80	0.94	3.68±0.77
BC	46.9	43.62±26.8	0.36	0.78	0.92	3.39±0.95
BCR	51.7	44.04±25.38	0.36	0.80	0.94	3.52±0.83
CTRL	50.0	24.48±11.98	0.40	0.84	0.93	3.63±0.75
BCR	56.0	–	–	–	–	3.61±0.69
WD	35.7	32.05±19.07	0.29	0.72	0.89	3.48±0.92
BCR	47.8	–	–	–	–	3.87±0.71

RESULTS

The below results correspond to the task of sentiment control:

- This indicates that some attributes are easier to control as compared to others
- We cannot clearly declare that CTRL or PPLM is superior

Method	Sentiment Acc. (%) (human)	Sentiment Acc. (%) (external classifier)	Perplexity (↓ better)	Dist-1 (↑ better)	Dist-2 (↑ better)	Dist-3 (↑ better)	Human Evaluation Fluency (↑ better)
B	19.3	52.2	42.1±33.14	0.37	0.75	0.86	3.54±1.08
BR	41.5	62.2	44.6±34.72	0.37	0.76	0.87	3.65±1.07
BC	39.6	64.4	41.8±34.87	0.33	0.70	0.86	2.79±1.17
BCR	73.7	78.8	46.6±40.24	0.36	0.77	0.91	3.29±1.07
CTRL	76.7	96.6	37.4±16.89	0.35	0.78	0.89	3.54±0.77
BCR	70.0	–	–	–	–	–	3.36±0.82
GPT2-FT-RL*	13.3	77.8	217.3±176.4	0.54	0.91	0.94	3.31±0.84
BCR	84.4	–	–	–	–	–	3.68±0.83
WD	18.9	52.2	31.7±28.0	0.33	0.69	0.83	3.67±0.89
BCR	61.1	–	–	–	–	–	3.75±0.66

RESULTS

[-] The chicken is now out on the grill. \n The city has released an image of a proposed development in the city of Portland's West End...

[Positive] The chicken was **delicious** – wonderfully moist, perfectly delicious, superbly fresh – and perfectly cooked. The only thing to say is that the sauce was **excellent**, and I think that the broth really complemented all of the other flavors. The **best** part was the sauce...

[Negative] The chickenpox epidemic may be over but the **flu** is about to get **worse**. The United States is facing one of the **worst** flu seasons on record and...

[-] The country's new chief minister, A.J. Paik, is a member of a group of prominent conservative politicians who have criticized the Obama administration's efforts to...

[Positive] The country's largest indoor painting event!\n Come **celebrate** with a **dazzling** display of **stunning** outdoor murals, a **stunning** display of art, and the world's **best** paint and art supplies from all over the world!

[Negative] The country's top **prison** system is forcing **prisoners** to use a **trash dump**, rather than a toilet, to flush their **waste** out, as the authorities **fear** the **waste** is more **toxic** and could cause **cancer**, an official at a major **prison** has revealed...



QUESTIONS?