Transfer Learning in Personalized Dialogue Generation

Haoyu Song

April 2020

A Simple Question

- What is the relation between **Pre-Training** and **Transfer Learning** ?

A Simple Question

- What is the relation between **Pre-Training** and **Transfer Learning** ?
  - **Pre-Training ⊆ Transfer Learning**
  - Transfer Learning has a wider scope. Details please refer to *A Survey on Transfer Learning*
  - In NLP, Pre-Training is the highlight of Transfer Learning
    - Word2Vec
    - BERT

Paper List

1. Neural Personalized Response Generation as Domain Adaptation, WWW Journal 2019

2. TransferTransfo-A Transfer Learning Approach for Neural Network Based Conversational Agents, AAAI 2019

3. Large-scale transfer learning for natural language generation, ACL 2019

4. A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data, AAAI 2020

Background

Rank	Creator	PPL	Hits@1	F1
1 🍐	🤗 (Hugging Face)	16.28🍎	80.7🍎	19.5🍎
2 🍐	ADAPT Centre	31.4	-	18.39
3 🍐	Happy Minions	29.01	-	16.01
4 🍐	High Five	-	65.9	-
5 🍐	Mohd Shadab Alam	29.94	13.8	16.91
6 🍐	Lost in Conversation	-	17.1	17.77
7 🍐	Little Baby(AI小奶娃)	-	64.8	-

2. AAAI 2019

3. ACL 2019

- NeurIPS 2018 Conversational Intelligence Challenge 2 (ConvAI2)
- http://convai.io/

Background

任务二：个性化对话竞赛排行榜

排名	更新日期	系统名称	机构名称	BLEU	Perplexity	Distinct
1	2019/7/14	persona-translator	**彩云科技&句子互动**	**0.0130**	**101.98**	**0.1689**
2	2019/7/5	fuxi persona	**网易伏羲实验室**	**0.0061**	**292.67**	**0.2160**
3	2019/7/6	Persona-dialogue	**中国科学院深圳先进技术研究院**	**0.0055**	**232.60**	**0.0520**
4	2019/7/15	wizare_smp	华南理工大学-CIKE实验室	0.0055	354.23	0.0872
5	2019/7/15	smp	WRFML LAB	0.0042	503.46	0.0467
6	2019/7/3	PAA_dialog	东北大学	0.0024	450.73	0.0447
7	2019/7/14	SIMPLE-Dialogue	NEU NLP LAB	0.0049	189.96	0.0325
8	2019/7/15	baseline1	复旦大学大数据学院	0.0009	1540.31	0.0437

4. AAAI 2020

- SMP 2019 The 3rd Evaluation of Chinese Human-Computer Dialogue Technology（SMP-ECDT 3）

Paper List

1. **Neural Personalized Response Generation as Domain Adaptation, WWW Journal 2019**

2. TransferTransfo-A Transfer Learning Approach for Neural Network Based Conversational Agents, AAAI 2019

3. Large-scale transfer learning for natural language generation, ACL 2019

4. A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data, AAAI 2020

1. Domain Adaptation-Information

- **Title:** Neural Personalized Response Generation as Domain Adaptation

- **Authors:** Wei-Nan Zhang, Ting Liu, Yifa Wang, Qingfu Zhu

- **Affiliation:** Harbin Institute of Technology

1. Domain Adaptation-Details

- **Goal:** To generate personalized dialogues based on Seq2Seq Model

Table 1: An example of the responses of different personality to a given post.

Post	Is it a proper dress for the first date?
Response #1	Yep.
Response #2	Honey, it is very suitable!
Response #3	It is better to wearing a silk scarf.

- No explicit persona texts are given

1. Domain Adaptation-Details

- **Models:** a two-phase approach, **initialization** and **adaptation**

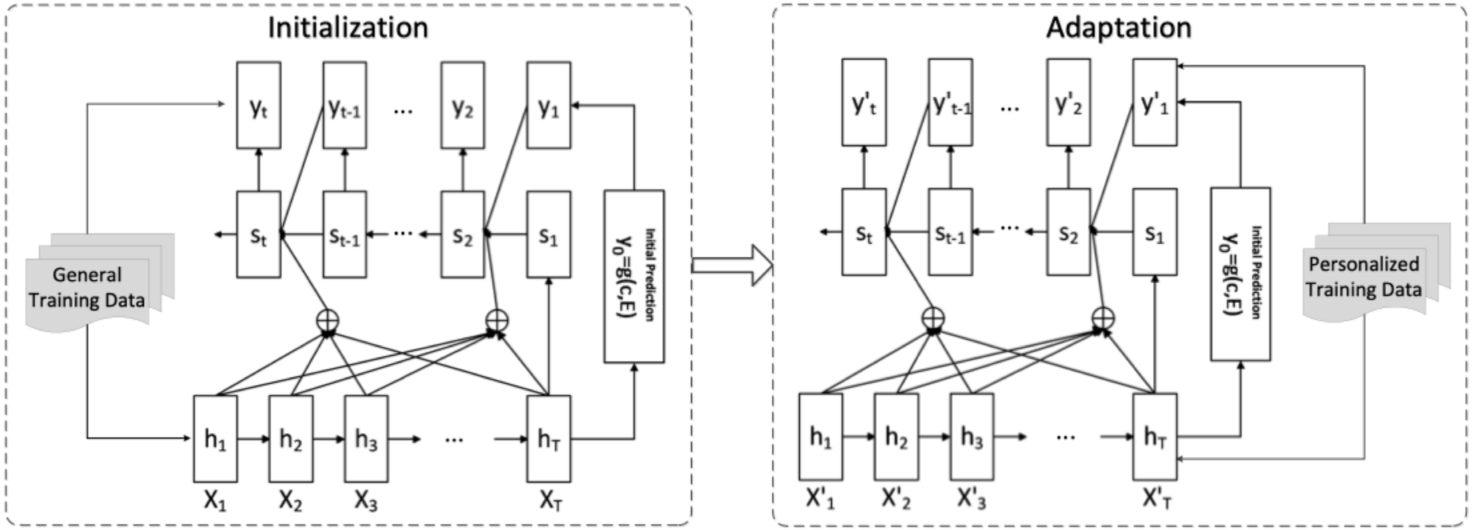

Figure 1: The framework of the proposed approach.

1. Domain Adaptation-Details

- **Dataset Collection**
  - General Data
    - 1 million one-to-one post-response pairs from several Chinese online forums, such as Weibo and Douban.
  - Personalized Data:
    - 5 volunteers, each shared 2,000 messages of their chatting history.
    - Then retrieve posts, which have similar responses to the personalized messages, from general data. The post-message pairs are the personalized data.

1. Domain Adaptation-Details

- **Experiment Results – imitation rate**

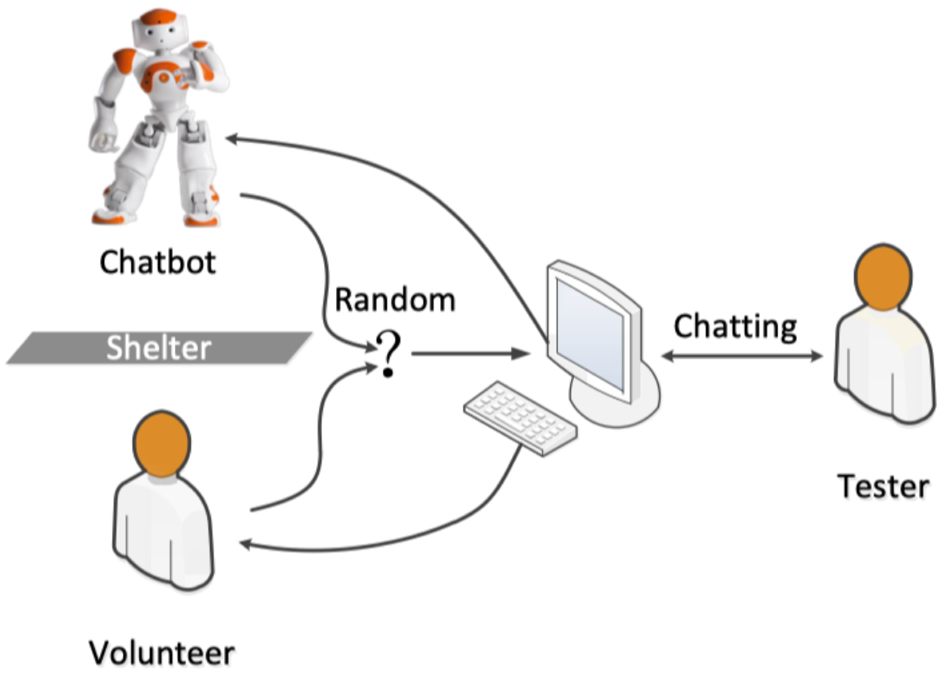

$$r_{imi} = \frac{n_{imi}}{n_{gr}}$$

1. Domain Adaptation-Details

- **Experiment Results – Human Evaluation**

Table 3: The experimental results of the proposed approach to personalized response generation. n_{gr} and n_{vr} represents the number of responses that are generated by the chatbot and the volunteer respectively. n_{test} is the total number of posts for testing. n_{imi} denotes the number of responses that is generated by the chatbot but are judged as the responses of the volunteer. r_{imi} denotes the imitation rate, which is defined in Equation (9).

	Volunteer #1	Volunteer #2	Volunteer #3	Volunteer #4	Volunteer #5	Sum
n_{gr}	29	26	21	33	33	142
n_{vr}	21	24	29	17	17	106
n_{test}	50	50	50	50	50	250
n_{imi}	11	9	8	13	9	50
r_{imi}	37.93%	34.62%	38.10%	39.40%	27.27%	35.21%

1. Domain Adaptation-Details

- **Experiment Results – Human Evaluation**

Table 4: The evaluation results of the 5 personalized response generation models by human judgement. n_{gr} represents the number of responses that are generated by the chatbot. n_{imi} denotes the number of responses that is generated by the chatbot but are judged as the responses of the volunteer. r_{imi} denotes the imitation rate, which is defined in Equation (9). PRM is short for the personalized responding model.

	PRM #1	PRM #2	PRM #3	PRM #4	PRM #5
n_{gr}	50	50	50	50	50
n_{imi}	6	8	8	13	10
r_{imi}	12%	16%	16%	26%	20%

1. Domain Adaptation-Details

- **Experiment Results – Word Statistics**

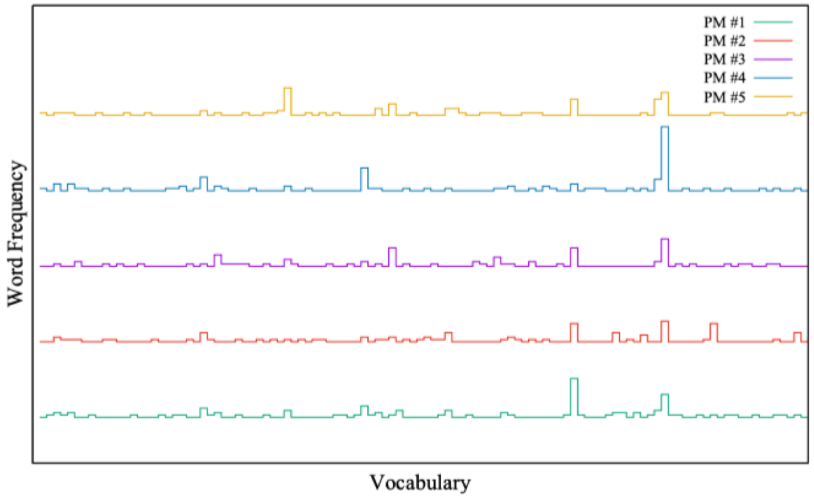

Figure 3: The distributions of the responding words on the vocabulary.

Table 6: The percentage of the overlapping words between the volunteer's responses and the chatbot generated responses for a same post set. "V" is short for "Volunteer".

	PRM #1	PRM #2	PRM #3	PRM #4	PRM #5
V #1	**20.00%**	17.86%	8.73%	14.78%	9.65%
V #2	15.45%	**18.64%**	8.13%	12.00%	10.83%
V #3	17.65%	10.75%	**22.22%**	14.77%	9.57%
V #4	17.63%	16.82%	8.28%	**17.92%**	11.32%
V #5	12.12%	12.96%	6.73%	10.28%	**18.52%**

1. Domain Adaptation-Details

- **Experiment Results – Some Cases**

Table 5: The real cases of the same post with different responses generated by PRM #1 - #5 models.

Post	Models	Responses
在哪能租到房子? (Where can I rent a house?)	PRM #1	我不知道那个网站。(I don't know the website.)
	PRM #2	去找一个吧。(To hunt a house.)
	PRM #3	市区。(The downtown.)
	PRM #4	联系中介吧。(You can contact the accommodation agency.)
	PRM #5	我这有地方住。(I have somewhere to stay.)
什么编程语言最好? (What is the best programming language?)	PRM #1	日语。(Japanese.)
	PRM #2	通常来说是Java。(Generally speaking, it is Java.)
	PRM #3	计算机语言。(The computer language.)
	PRM #4	不用。(I don't need.)
	PRM #5	PHP。(PHP.)
洗衣服好累啊。 (It is too tired to wash the clothes.)	PRM #1	用洗衣机洗。(Use the washing machine.)
	PRM #2	累。(I'm tired.)
	PRM #3	我能。(I can.)
	PRM #4	恩。(Yes.)
	PRM #5	去休息吧。(Have a rest.)

1. Domain Adaptation-Summaries

- This work was initialized at the end of 2016. At that time, transfer learning and pre-training in dialogues are not as ubiquitous as today.

- This work is suitable as a baseline:

  - It is a published paper

  - Easy to implement

  - Easy to outperform

Paper List

1. Neural Personalized Response Generation as Domain Adaptation, WWW Journal 2019
2. **TransferTransfo-A Transfer Learning Approach for Neural Network Based Conversational Agents, AAAI 2019**
3. Large-scale transfer learning for natural language generation, ACL 2019
4. A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data, AAAI 2020

2. TransferTransfo-Information

- **Title:** TransferTransfo: A Transfer Learning Approach for Neural Network Based Conversational Agents

- **Authors:** Thomas Wolf, Victor Sanh, Julien Chaumond and Clement Delangue

- **Affiliation:** HuggingFace Inc.

2. TransferTransfo-Motivation

- **The well-known challenges in open-domain dialogues:**
  - The lack of a consistent personality
  - The absence of a long-term memory
  - A tendency to produce consensual and generic responses

2. TransferTransfo-Details

- **Goals**
  - To generate persona-based responses (a generative model)
  - To select 1 response from 20 candidates (a retrieval model)
- **Metrics**
  - Perplexity, F1, Hits@1
  - Real-time human interactive evaluation

2. TransferTransfo-Details

- **Model**
  - TransferTransfo → Transfer Transformer
  - A 12-layer decoder-only transformer with masked self-attention heads (768 dimensional states and 12 attention heads)
  - It is in the same architecture to the openAI GPT. But the GPT 2018 did not prove its effectiveness in generation tasks (only verified on NLU tasks).

2. TransferTransfo-Details

- **Dataset**

Persona 1	Persona 2
I like to ski	I am an artist
My wife does not like me anymore	I have four children
I have went to Mexico 4 times this year	I recently got a cat
I hate Mexican food	I enjoy walking for exercise
I like to eat cheetos	I love watching Game of Thrones

[PERSON 1:] Hi
[PERSON 2:] Hello ! How are you today ?
[PERSON 1:] I am good thank you , how are you.
[PERSON 2:] Great, thanks ! My children and I were just about to watch Game of Thrones.
[PERSON 1:] Nice ! How old are your children?
[PERSON 2:] I have four that range in age from 10 to 21. You?
[PERSON 1:] I do not have children at the moment.
[PERSON 2:] That just means you get to keep all the popcorn for yourself.
[PERSON 1:] And Cheetos at the moment!
[PERSON 2:] Good choice. Do you watch Game of Thrones?
[PERSON 1:] No, I do not have much time for TV.
[PERSON 2:] I usually spend my time painting: but, I love the show.

2. TransferTransfo-Details

- **Dataset**

Persona 1	Persona 2
I like to ski	I am an artist
My wife does not like me anymore	I have four children
I have went to Mexico 4 times this year	I recently got a cat
I hate Mexican food	I enjoy walking for exercise
I like to eat cheetos	I love watching Game of Thrones

[PERSON 1:] Hi
[PERSON 2:] Hello ! How are you today ?
[PERSON 1:] I am good thank you , how are you.
[PERSON 2:] Great, thanks ! My children and I were just about to watch Game of Thrones.
[PERSON 1:] Nice ! How old are your children?
[PERSON 2:] I have four that range in age from 10 to 21. You?
[PERSON 1:] I do not have children at the moment.
[PERSON 2:] That just means you get to keep all the popcorn for yourself.
[PERSON 1:] And Cheetos at the moment!
[PERSON 2:] Good choice. Do you watch Game of Thrones?
[PERSON 1:] No, I do not have much time for TV.
[PERSON 2:] I usually spend my time painting: but, I love the show.

2. TransferTransfo-Details

- **Pre-training Data**
  - The BooksCorpus dataset (Zhu et al., 2015), containing over 7,000 unpublished books (about 800M words) from a variety of genres.
  - Using the document-level corpus rather than a shuffled sentence-level corpus.
  - Taking advantage of long contiguous sequences and paragraphs and learn to condition on long-range information.

2. TransferTransfo-Details

- **Input Representation**

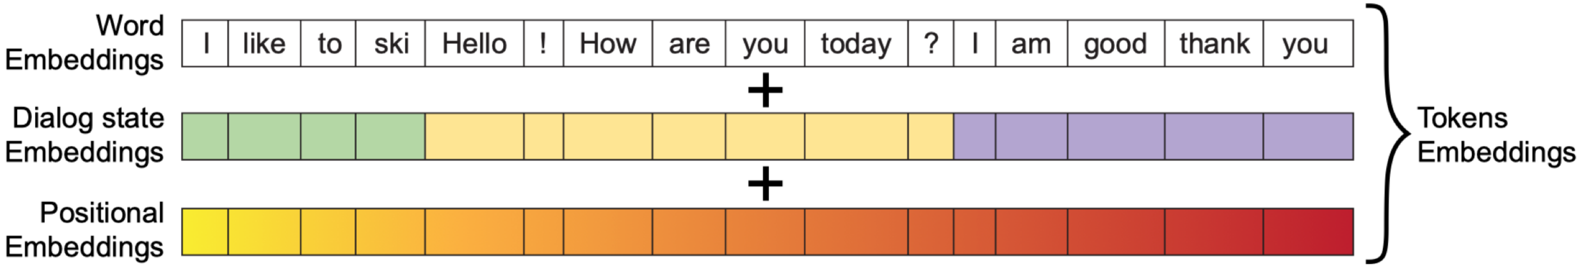

Figure 1: TranferTransfo's input representation. Each token embedding is the sum of a word embedding, a dialog state embedding and a positional embedding.

- input sequence = persona texts + dialogue history
- dialogue history = speak_A_1 + speak_B_1 + speak_A_2 + ⋯
- reusing the same positional embeddings for each persona texts to promote an invariance to persona texts ordering

2. TransferTransfo-Details

- **Multi-task learning**
  - A language model loss (cross-entropy)
  - A next-utterance classification loss

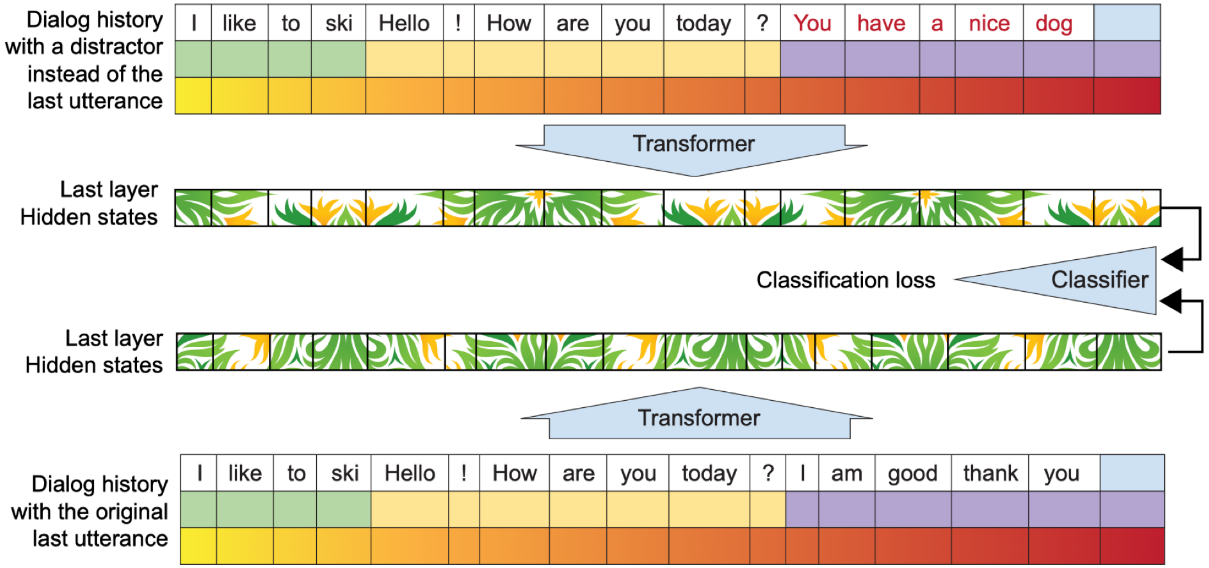

Similar to the Next Sentence Prediction task in BERT (a parallel work):

Training a head to distinguish a correct next utterance appended to the input sequence from a set of randomly sampled distractors (in practice between 2 and 6 randomly sampled utterances).

2. TransferTransfo-Details

- **Fine-tuning details**
  - batch size of 32, an average of 250 tokens
  - 200,000 steps, about 2 epochs on PersonaChat
  - Adam with a learning rate of 6.25e-5, $\beta 1 = 0.9$, $\beta 2 = 0.999$
  - dropout probability of 0.1 on all layers
  - 10 hours on four K80 GPUs

2. TransferTransfo-Details

- **Results**

	Eval			Test		
Model	PPL	Hits@1	F1	PPL	Hits@1	F1
Generative Profile Memory (Zhang et al., 2018)	*34.54*	12.5	–	–	–	–
Retrieval KV Profile Memory (Zhang et al., 2018)	–	*51.1*				
Seq2Seq + Attention (ConvAI2 baseline[3])	35.07	12.5	*16.82*	*29.8*	12.6	*16.18*
Language Model (ConvAI2 baseline[4])	51.1	–	15.31	46.0	–	15.02
KV Profile Memory (ConvAI2 baseline[5])	–	55.1	11.72	–	*55.2*	11.9
TransferTransfo (this work)	**17.51**	**82.1**	**19.09**	**16.28**	**80.7**	**19.5**

- **Hits@1 of two retrieval models (on the hidden Test set):**

4	High Five	-	65.9	-
7	Little Baby(AI小奶娃)	-	64.8	-

2. TransferTransfo-Details

- **Are these metrics good ? (excerpt from competition summary)**

LESSONS?

- ***How good are these automated metrics?***
  - There was **some** correlation between PPL and hits@1and human evaluation

- **However, Major ISSUE with F1:**
  - Mega-dumb baseline pick a combination of frequent words:
    - "**i am you to do the a, and your is like!?**"

  each turn would give the **best F1 score in the competition** 😂

  (**19.6** on the test set and **20.5** on the valid set compared to Hugging Face's 19.5 and 19.1)

- In any case, shown not be correlated well with human metrics (Liu et al, 2016)

2. TransferTransfo-Summaries

- This paper presented their work in ConvAI2, and its superior performance largely benefits from the pretrained GPT.

- This work is an early attempt in leveraging pre-trained LM for dialogue generation. Before BERT (Oct. 2018) and GPT (Jun. 2018), few people knew the power of pretrained transformers.

- Although it is a short paper, it has been cited by 38 times since 2019 and has an impact on later work such as DialogueGPT.

Paper List

1. Neural Personalized Response Generation as Domain Adaptation, WWW Journal 2019

2. TransferTransfo-A Transfer Learning Approach for Neural Network Based Conversational Agents, AAAI 2019

3. **Large-scale transfer learning for natural language generation, ACL 2019**

4. A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data, AAAI 2020

3. Lost in Conversation-Information

- **Title:** Large-Scale Transfer Learning for Natural Language Generation

- **Authors:** Sergey Golovanov, Rauf Kurbanov, Sergey Nikolenko et al.

- **Affiliation:** Neuromation OU, Estonia; Steklov Mathematical Institute at St.Petersburg.

3. Lost in Conversation-Details

- **Model**

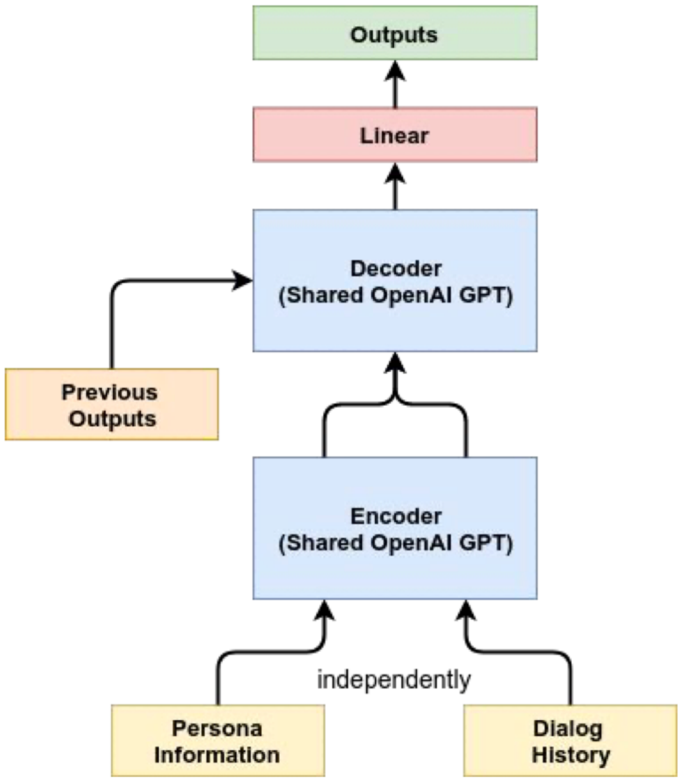

- Shared encoder and decoder - pretrained OpenAI GPT
- Shared pre-softmax linear layer and token embeddings
- Beam-search with length penalty and annealing for improving answer diversity
- Reduction of persona information and dialog history – first and last 512 tokens respectively

3. Lost in Conversation-Details

- **Model**

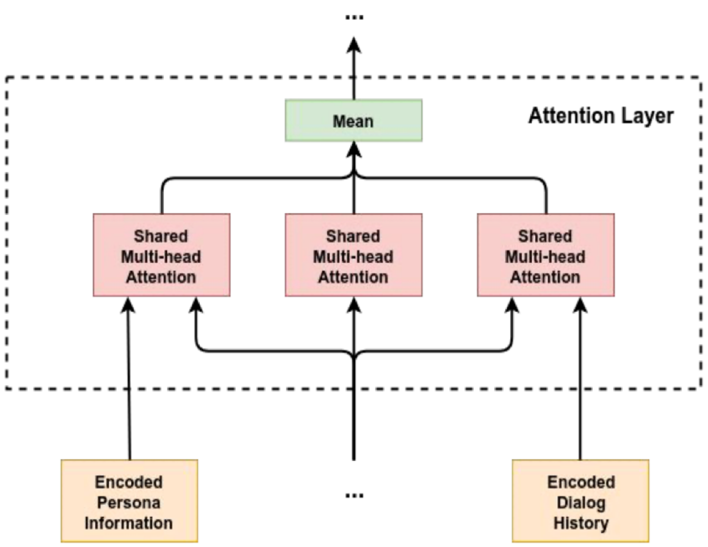

Attention layer modifications:

- Shared multi-head attention layers
- Parallel computation of attention for inputs
- Merge of attentions - mean

3. Lost in Conversation-Details

- **Dataset:** PersonaChat
- **Pre-Training Dataset**
  - DailyDialog
    - Li Y. et al. DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset
  - Reddit comments dataset
    - files.pushshift.io/reddit/comments

3. Lost in Conversation-Details

- **Training Settings**
  - batch size: 256
  - learning rate: 6.25e-5
  - warmup: 16000
  - label smoothing: 0.1
  - dropout: 0.1

  - First stage: one week on Nvidia GTX 1080TI
  - Finetuning stage: about two days on Nvidia GTX 1080TI

3. Lost in Conversation-Details

- **The Interesting Part of This Paper: Compare two architectures**

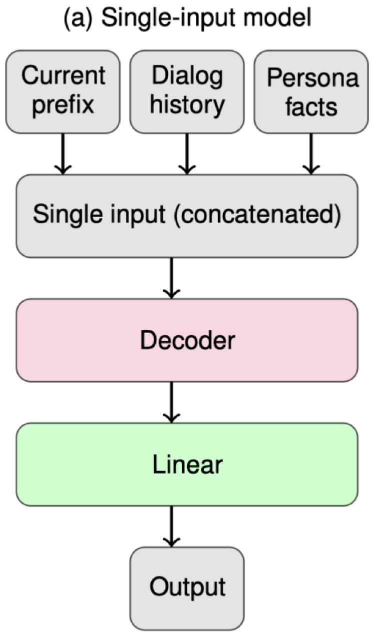

Decoder-Only

VS

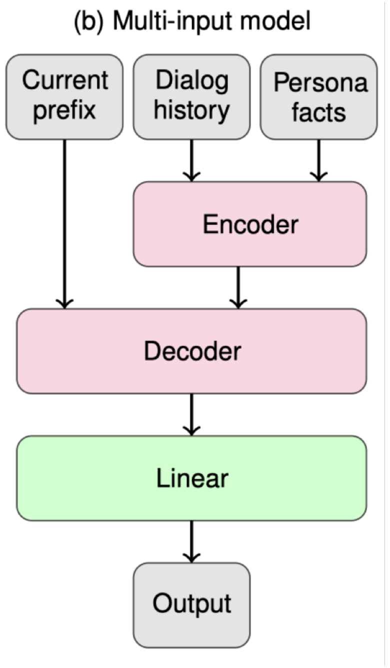

Encoder-Decoder

3. Lost in Conversation-Details

- **The Interesting Part of This Paper: Compare Two Architectures**
  - Both architectures reach comparable performances on the automatic metrics

Model	METEOR	NIST-4	BLEU	Entropy-4	Distinct-2	Average Length
Single-input (zero-shot)	0.07727	1.264	2.5362	9.454	0.1759	9.671
Single-input (additional embeddings)	0.07641	1.222	2.5615	9.234	0.1614	9.43
Multi-input	**0.07878**	**1.278**	**2.7745**	9.211	0.1546	9.298

3. Lost in Conversation-Details

- **The Interesting Part of This Paper: Compare Two Architectures**

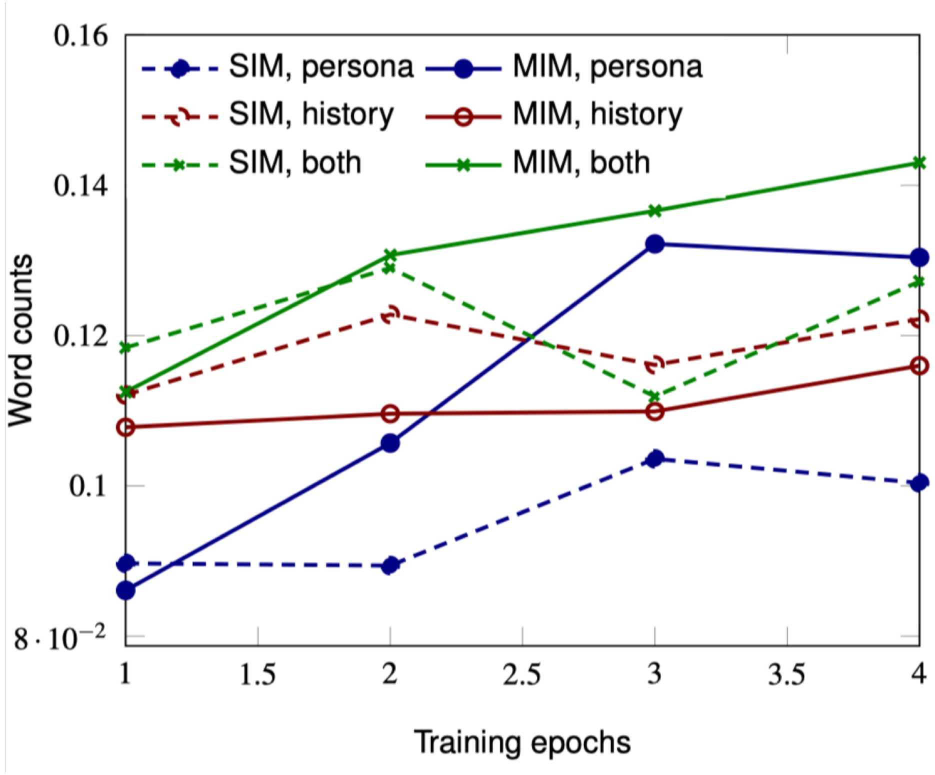

Words are labeled to three categories:
1. words that were mentioned in the **persona texts**;
2. words that were mentioned in the **dialog history**;
3. words that were mentioned in **both**

Conclusion:
1. Single-input stays closer to the dialog history
2. Multi-input stays closer to persona texts

Possible Reasons:
Persona texts are **not ordered**. In Single-input model, they are handled sequentially. Older history becomes less relevant.

3. Lost in Conversation-Summaries

- This model got the best human evaluation results in ConvAI2 final round.

Human Evaluation Leaderboard

Rank	Creator	Rating
1 🍐 🍌	Lost in Conversation [code]	3.11🍌
2 🍐 🍎 🍎 🍎	🤗 (Hugging Face)	2.68
3 🍐	Little Baby(AI小奶娃)	2.44

- It is an educated guess that encoder-decoder structure is more suitable for dialogue generation tasks.

Paper List

1. Neural Personalized Response Generation as Domain Adaptation, WWW Journal 2019

2. TransferTransfo-A Transfer Learning Approach for Neural Network Based Conversational Agents, AAAI 2019

3. Large-scale transfer learning for natural language generation, ACL 2019

4. **A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data, AAAI 2020**

4. Pre-training Persona-sparse Generation

- **Title:** A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data

- **Authors:** Yinhe Zheng, Rongsheng Zhang, Xiaoxi Mao, Minlie Huang

- **Affiliation:** Tsinghua University; NetEase Inc.; Samsung Research

4. Pre-training Persona-sparse Generation

- **What is "persona-sparse"?**
  - Most speakers in daily conversations are not aiming to exhibit their personas within limited turns of interactions.
  - Data collected from real-world conversations only contain a limited amount of dialogues that relate to speakers' persona.
  - In contrast, **PersonaChat** is a **persona-dense** dataset. Its collection scheme guaranteed to yield dialogues that cover rich persona features.

4. Pre-training Persona-sparse Generation

- **Dataset**

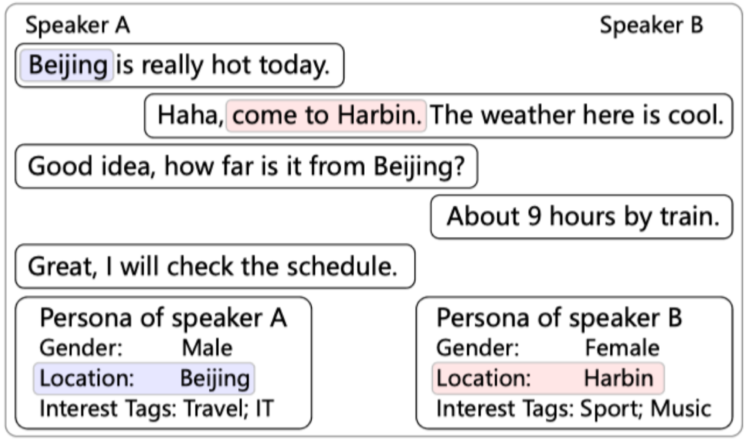

Total number of dialogues	5.44 M
Total number of speakers	1.31 M
Total number of utterances	14.40 M
Dialogues with more than 4 utterances	0.81 M
Average utterances per dialogue	2.65
Average tokens per utterance	9.46

- SMP-ECDT 2019 Dataset
- Collected from Chinese social media Weibo
- Weibo post and its replies, together with a structured profile of each speaker
- **Random test set** contains 10K sessions of randomly sampled dialogues
- **Biased test set** provides contexts which speakers tend to reveal their personas, selected by humans.

4. Pre-training Persona-sparse Generation

- **Pre-training Dataset**
  - A dataset collected from a set of Chinese novels, which covered a variety of genres (including Comedy, Romance, Mystery).
  - The final pre-training corpus contains about 0.5 billion tokens.
  - a character-level language model with a vocabulary size of 13,084.

4. Pre-training Persona-sparse Generation

- **Solutions**
  - a pre-training based method that can utilize persona-sparse data
  - an **attention routing mechanism** to weigh persona features dynamically in the decoder

4. Pre-training Persona-sparse Generation

- **Encoding Persona**

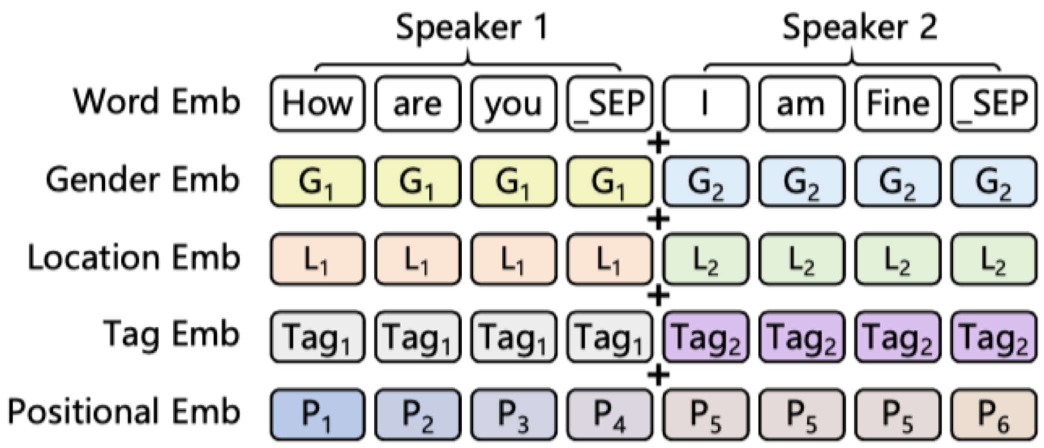

- Sum of a word embedding, a positional embedding and attribute embeddings
- Attribute embedding is obtained by utilizing look-up tables
- Interest Tags is computed as the average of all the tag embeddings

4. Pre-training Persona-sparse Generation

- **Model**

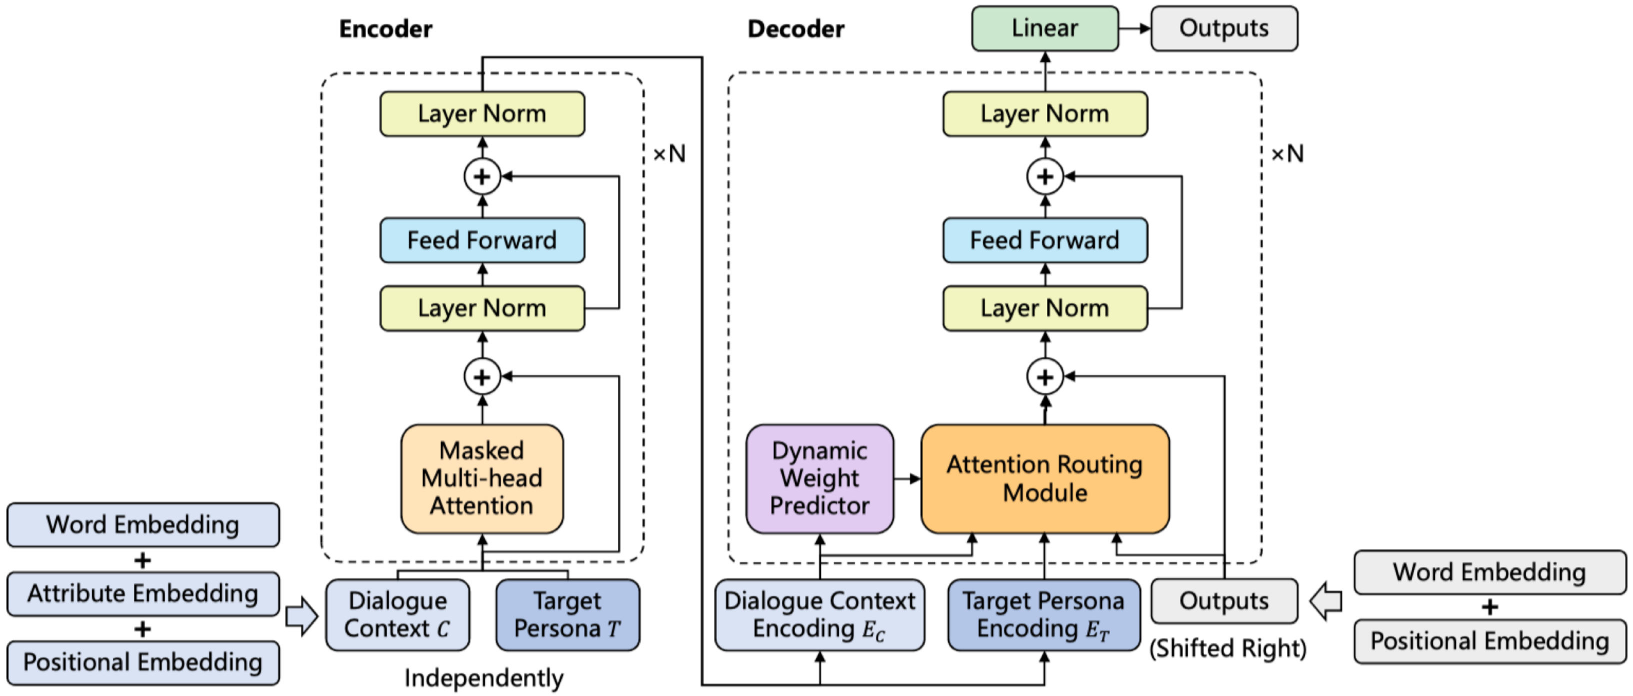

Figure 2: The framework of the proposed personalized dialogue generation model. The encoder and decoder share the same set of parameters. The dialogue context and the target persona are encoded independently using the encoder and their encodings are fed into the attention routing module in each decoder block. A dynamic weight predictor is trained to weigh the contribution of each route.

4. Pre-training Persona-sparse Generation

- **Attention Routing**

$$O_T = \mathrm{MultiHead}(E_{prev}, E_T, E_T) \quad (2)$$
$$O_C = \mathrm{MultiHead}(E_{prev}, E_C, E_C) \quad (3)$$
$$O_{prev} = \mathrm{MultiHead}(E_{prev}, E_{prev}, E_{prev}) \quad (4)$$

$$O_{merge} = \alpha O_T + (1-\alpha) O_C + O_C + O_{prev} \quad (5)$$

- E_{prev} previous decoded words
- E_T target profile
- E_C dialogue context
- α the confidence of the response is persona related

4. Pre-training Persona-sparse Generation

- **Automatic Results**

Table 2: Automatic evaluation on the random test set.

Model	Acc.	BLEU	F1	Dist.	ppl.
Att+PAB	13.99	1.61	8.60	0.130	69.30
Trans.	7.80	3.97	12.51	0.132	43.12
TTransfo	8.80	4.06	12.63	0.169	**32.12**
TTransfo+P	43.05	3.44	11.28	0.158	43.78
LConv	9.45	4.19	12.99	0.157	32.64
LConv+P	48.00	3.56	11.46	0.136	42.00
Ours	32.80	4.18	12.52	**0.171**	35.06
Ours, α=1	**84.55**	3.45	10.96	0.154	38.56
Ours, α=0	12.90	**4.56**	**13.02**	**0.171**	33.71
w/o PreT	27.10	3.86	11.62	0.146	48.48
w/o AEmb	31.85	4.15	12.56	0.164	35.75
w/o DWP	30.70	4.15	12.34	0.169	34.10
+ HW	32.55	3.50	11.90	0.151	38.52

Table 3: Automatic evaluation on the biased test set.

Model	Acc.	BLEU	F1	Dist.	ppl.
Att. + PAB	47.60	3.08	12.50	0.133	94.38
Trans.	34.93	7.06	15.38	0.203	85.80
TTransfo	45.87	8.68	17.39	0.260	**34.83**
TTransfo+P	61.61	9.10	18.41	0.257	38.07
LConv	44.34	8.47	17.08	0.238	37.44
LConv+P	59.88	9.82	18.91	0.231	41.68
Ours	92.13	10.53	19.47	0.256	38.68
Ours, α=1	**94.24**	**11.63**	**20.51**	**0.262**	39.74
Ours, α=0	51.44	9.00	17.44	0.249	40.89
w/o PreT	71.74	9.36	18.29	0.222	95.00
w/o AEmb	73.51	10.51	19.41	0.247	39.36
w/o DWP	73.90	10.61	19.26	0.256	37.08
+ HW	69.87	9.01	19.81	0.232	36.37

4. Pre-training Persona-sparse Generation

- **Effect of α**

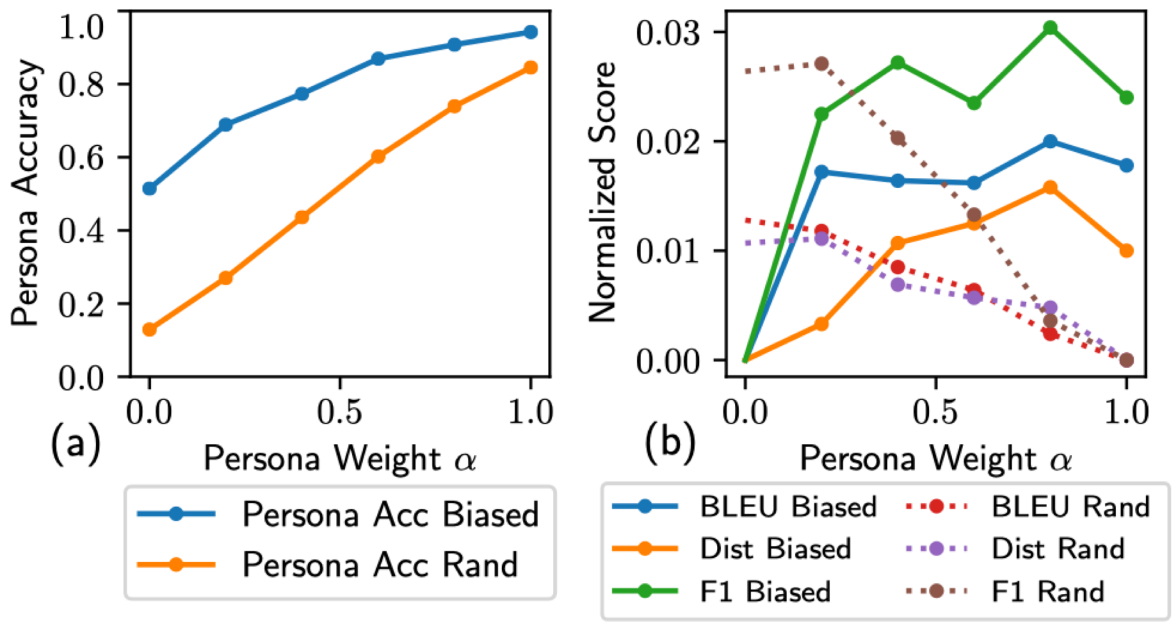

- α controls the persona information
- In the persona-sparse scenario, too many persona information could lead to the decrease of performance

4. Pre-training Persona-sparse Generation

- **Human Evaluations and Cases**

Model	Utterance Fluency		Persona Consistency		Context Coherency	
	Rand	Biased	Rand	Biased	Rand	Biased
Trans.	1.852	1.810†	0.997†	1.068†	1.428†	1.500
TTransfo	1.832†	1.890	1.015†	1.100†	1.498	1.517
TTransfo+P	1.802†	1.837†	1.125†	1.195†	1.217†	1.483†
LConv	1.863	1.882	1.028†	1.147†	1.490	1.550
LConv+P	1.832†	1.875†	1.093†	1.173†	1.238†	1.478†
Ours	1.837†	**1.912**	1.092†	1.198†	1.487	**1.563**
Ours, α=1	1.835†	1.900	**1.248**	**1.268**	1.303†	1.467†
Ours, α=0	**1.890**	1.880†	0.997†	1.085†	**1.535**	1.463†
Gold Resp	1.928	1.922	1.015	1.423	1.758	1.807

† significant difference with the best result (t-test, p-value<0.05)

Dialogue Context:	I am free this weekend. It that OK? (我这周末闲，好不) (Gender: Female, Location: Nanjing, Interest Tags: Cate)
Target Persona:	(Gender: Male, Location: Nanjing, Interest Tags: Null)
Trans. :	I'm also free, but I don't know what's going on. (我也闲，但是我不知道怎么回事)
TTransfo :	Great, come here and we can have a meal together. (好呀，你来了我们一起吃个饭呗)
TTransfo+P :	I am free this weekend. It that Okay? (我这周末闲，好不好？)
LConv :	It is not OK. (不好)
LConv+P :	Are you in Nanjing or Nanjing? (你在南京还是南京啊？)
Ours :	Ok, are you coming to Nanjing? (好啊，你来南京吗？)
Ours, $\alpha = 1$:	Come to Nanjing, I will treat you to dinner and film. (来南京，我请你吃饭看电影)
Ours, $\alpha = 0$:	Ok, where are you? I will go back tomorrow. (好啊，你在那呢？我明天就回去了)
Gold Resp :	OK, come to Nanjing! (好，来南京！)

* with scale 0, 1, 2

Conclusion

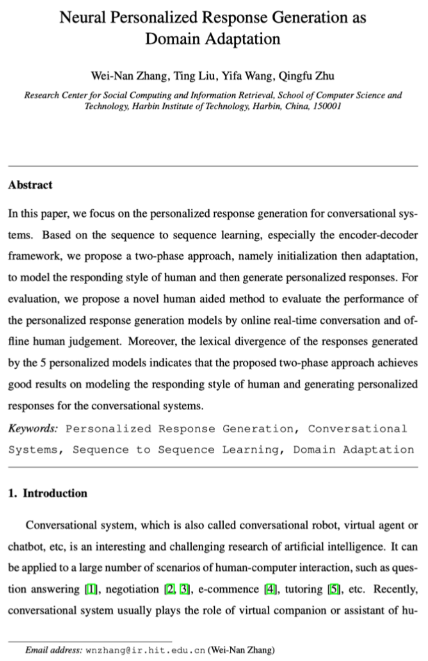

Neural Personalized Response Generation as Domain Adaptation

Wei-Nan Zhang, Ting Liu, Yifa Wang, Qingfu Zhu

Research Center for Social Computing and Information Retrieval, School of Computer Science and Technology, Harbin Institute of Technology, Harbin, China, 150001

Abstract

In this paper, we focus on the personalized response generation for conversational systems. Based on the sequence to sequence learning, especially the encoder-decoder framework, we propose a two-phase approach, namely initialization then adaptation, to model the responding style of human and then generate personalized responses. For evaluation, we propose a novel human aided method to evaluate the performance of the personalized response generation models by online real-time conversation and offline human judgement. Moreover, the lexical divergence of the responses generated by the 5 personalized models indicates that the proposed two-phase approach achieves good results on modeling the responding style of human and generating personalized responses for the conversational systems.

Keywords: Personalized Response Generation, Conversational Systems, Sequence to Sequence Learning, Domain Adaptation

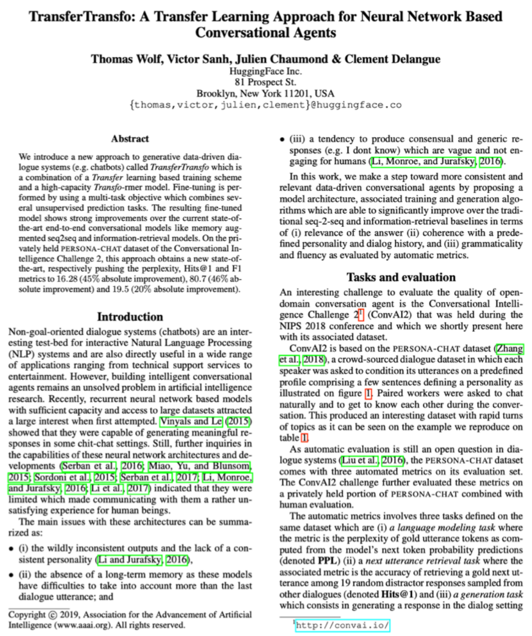

TransferTransfo: A Transfer Learning Approach for Neural Network Based Conversational Agents

Thomas Wolf, Victor Sanh, Julien Chaumond & Clement Delangue

HuggingFace Inc.

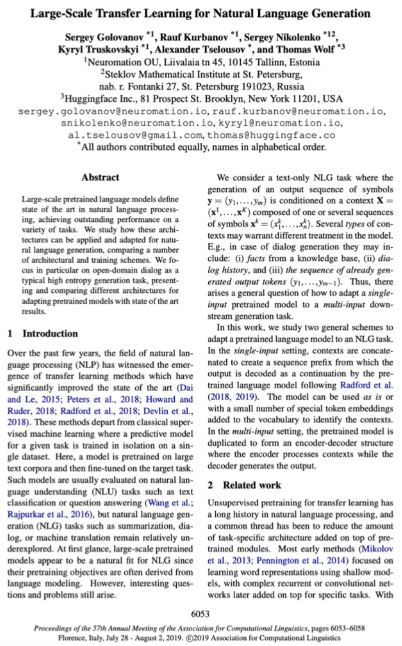

Large-Scale Transfer Learning for Natural Language Generation

A Pre-training Based Personalized Dialogue Generation Model with Persona-sparse Data

Yinhe Zheng, Rongsheng Zhang, Xiaoxi Mao, Minlie Huang

- Domain adaption
- No explicit persona
- RNN-based model
- 2016-2017

- Pre-training + Finetuning
- With explicit persona
- Transformer-based model
- 2018-2020

SUMMARY

1. **Better language model, better dialogue generation model**
2. **Transfer learning in personalized dialogue is still under-explored**

END
Q & A