北京大學
PEKING UNIVERSITY

Personalizing Adaptation for Meta-Learning-based Dialogue Generation

Zequn Liu, Bin Feng

Overview

- Personalized Dialogue Generation
- Related Work
- Model
- Experiments

Personalized Dialogue Generation

- Response generation depends on the persona
  - Q: Do you have any pets?
  - P1: He has a dog.   A1: Yes, I have a dog
  - P2: She doesn't have a pet. A2: No.
- Difficulties
  - The response should be consistent with the persona   Personality
  - Each persona doesn't have enough training data   Quality

Related Work

- Use the persona descriptions to generate response
  - Use all the users to train a model
  - Input: query + persona description
  - Eg： Zhang et al.(2018) propose to calculate the attention of the current query over all the description sentences, then use this attention to re-write the generated replies.
  - However, persona description are often unavailable.
- Meta-learning-based dialogue generation
  - Train a model for each persona
  - Do not need the persona description

Saizheng Zhang, Emily Dinan, Jack Urbanek, ArthurSzlam, Douwe Kiela, and Jason Weston. Personalizing dialogue agents: I have a dog, do you have pets too? In ACL, pages 2204–2213, 2018

Related Work

- Meta-learning-based Dialogue Generation

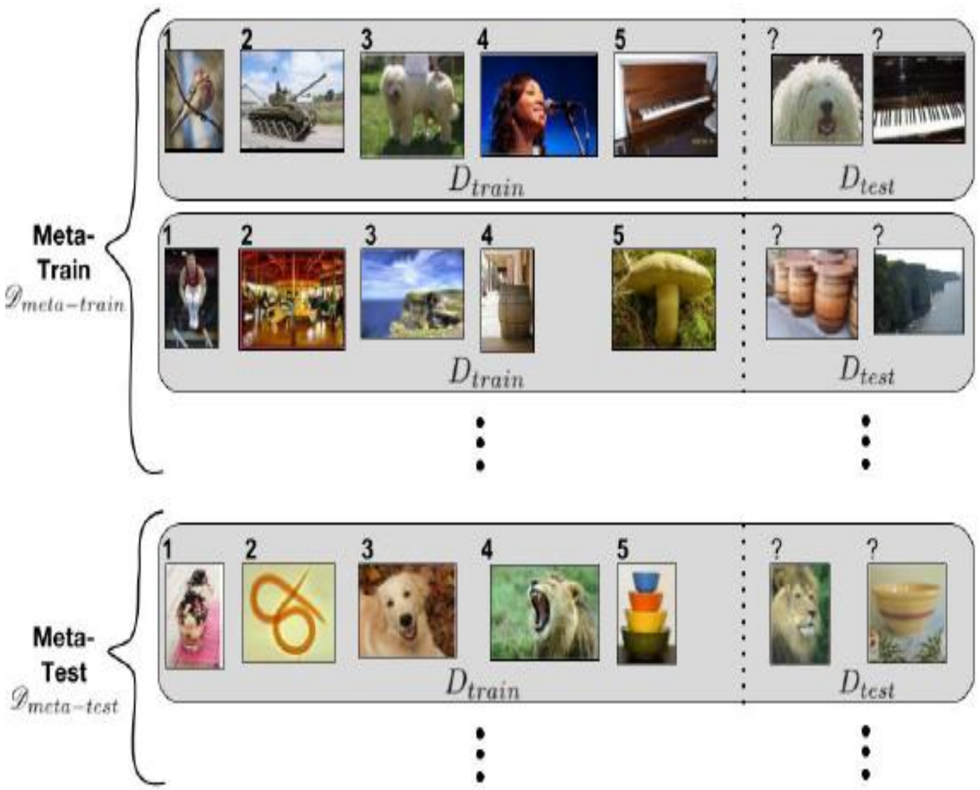

Meta-learning:

1. Train a meta-learner
2. Fast adaptaion

D_{train}, D_{test}

D_{train}, D_{test}

D_{train}, D_{test}

Meta Train

D_{train}, D_{test}

D_{train}, D_{test}

MAML

- Meta-learner: Find an initialization for fast adaptation

Algorithm 2 MAML for Few-Shot Supervised Learning

Require: $p(\mathcal{T})$: distribution over tasks
Require: α, β: step size hyperparameters

randomly initialize θ
while not done **do**
  Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
  for all $\mathcal{T}_i$ **do**
    Sample K datapoints $\mathcal{D} = \{\mathbf{x}^{(j)}, \mathbf{y}^{(j)}\}$ from $\mathcal{T}_i$
    Evaluate $\nabla_\theta \mathcal{L}_{\mathcal{T}_i}(f_\theta)$ using $\mathcal{D}$ and $\mathcal{L}_{\mathcal{T}_i}$ in Equation (2) or (3)
    Compute adapted parameters with gradient descent: $\theta_i' = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_i}(f_\theta)$
    Sample datapoints $\mathcal{D}_i' = \{\mathbf{x}^{(j)}, \mathbf{y}^{(j)}\}$ from $\mathcal{T}_i$ for the meta-update
  end for
  Update $\theta \leftarrow \theta - \beta \nabla_\theta \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta_i'})$ using each $\mathcal{D}_i'$ and $\mathcal{L}_{\mathcal{T}_i}$ in Equation 2 or 3
end while

MAML in personalized dialogue generation

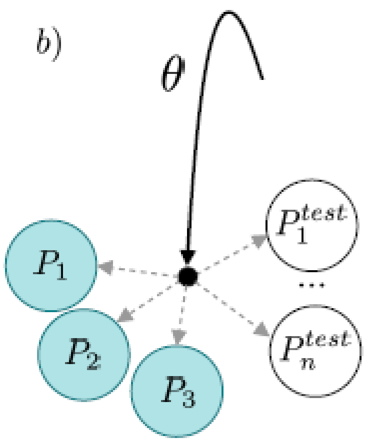

Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung.Personalizing dialogue agents via meta-learning. In ACL, pages 5454–5459, 2019.

Our Model

- MAML-based Personalized Dialogue Generation
  - Hyperparameters of the adaptation process is the same for different tasks
    - Learning rate, steps,…
    - Different tasks need different hyperparameters in personalized dialogue generation (Liu et al. 2020)
- Use a neural network to generate hyperparameters for each task
- Notation

$$\mathscr{D} = \{\mathcal{D}_{p_1}, \ldots, \mathcal{D}_{p_z}\} \qquad \mathcal{D}_p = \{U_1, \ldots, U_k\}.$$

$$X = \{u_1, \ldots, u_{t-1}\}. \qquad Y = u_t$$

Zequn Liu, Ruiyi Zhang, Yiping Song, and Ming Zhang. When does maml work the best? an empirical study on model-agnostic meta-learning in nlp applications, arXiv preprint arXiv:2005.11700. 2020.

Our Model

- Dialogue Generation Model: Seq2seq
- Adaptation on Encoder
  - The same as MAML

$$\theta^{(1)}_{p_i,e} = \theta_{0,e} - \alpha \nabla_{\theta_e} \mathcal{L}_{\mathcal{D}^{train}_{p_i}}(\theta_0)$$
$$\theta^{(t)}_{p_i,e} = \theta^{(t-1)}_{p_i,e} - \alpha \nabla_{\theta_e} \mathcal{L}_{\mathcal{D}^{train}_{p_i}}\left(\theta^{(t-1)}_{p_i}\right) \quad (1)$$

- Adaptation on Decoder
  - Meta-LSTM (Ravi et al. 2016)

$$\theta^{(t)}_{p_i,d} = f^{(t)}_{p_i} \odot \theta^{(t-1)}_{p_i,d} - i^{(t)}_{p_i} \odot \nabla_{\theta_d} \mathcal{L}_{\mathcal{D}^{train}_{p_i}}\left(\theta^{(t-1)}_{p_i}\right)$$
$$f^{(t)}_{p_i} = \sigma(W_f \cdot [\nabla_{\theta_d} \mathcal{L}_{\mathcal{D}^{train}_{p_i}}\left(\theta^{(t-1)}_{p_i}\right), \theta^{(t-1)}_{p_i,d}, \mathcal{L}_{\mathcal{D}^{train}_{p_i}}\left(\theta^{(t-1)}_{p_i}\right), f^{(t-1)}_{p_i}] + b_f)$$
$$i^{(t)}_{p_i} = \sigma(W_i \cdot [\nabla_{\theta_d} \mathcal{L}_{\mathcal{D}^{train}_{p_i}}\left(\theta^{(t-1)}_{p_i}\right), \theta^{(t-1)}_{p_i,d}, \mathcal{L}_{\mathcal{D}^{train}_{p_i}}\left(\theta^{(t-1)}_{p_i}\right), i^{(t-1)}_{p_i}] + b_i)$$
$$\theta^{(0)}_{p_i,d} = c_0 \quad (3)$$

scale

Learning rate

Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016

Our Model

- Training
  - Meta-objective $\min_{\theta_{0,e},\theta_m} \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}\left(\theta_{p_i}^{(T)}\right)$

$$\theta_m = [W_i, W_f, b_i, b_f, c_0].$$

$$\theta_{0,e} \leftarrow \theta_{0,e} - \beta \nabla_{\theta_{0,e}} \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}\left(\theta_{p_i}^{(T)}\right)$$
$$\theta_m \leftarrow \theta_m - \beta \nabla_{\theta_m} \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}\left(\theta_{p_i}^{(T)}\right) \quad (5)$$

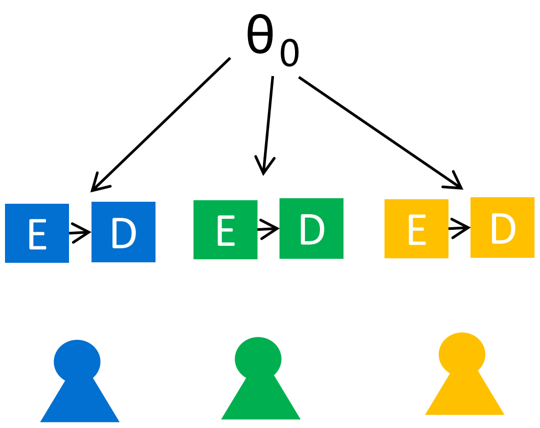

MAML

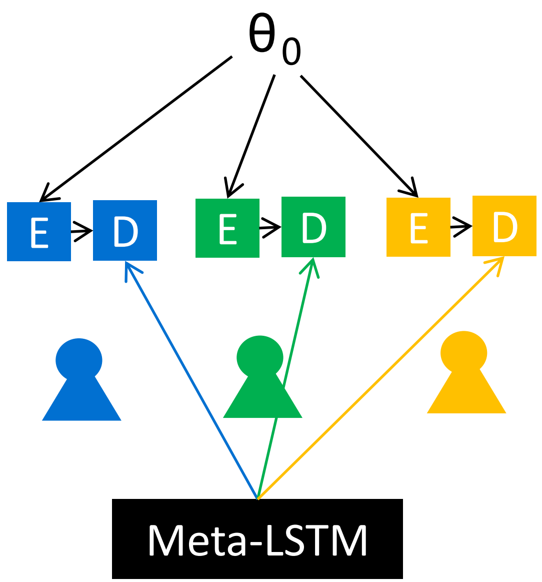

Our Model

Algorithm 1: Proposed Training Algorithm

Input: α, β: Learning rate
$\mathcal{D}_{train}$: meta-training data
randomly initialize $\theta_{0,e}, \theta_m$.
while *not done* **do**
  Sample a persona $\mathcal{D}_{p_i} \sim \mathscr{D}_{train}$
  for *t=1,T* **do**
    Adapt $\theta_{p_i,e}$ using Equation 1
    Adapt $\theta_{p_i,d}$ using Equation 3
  update $\theta_{0,e}, \theta_m$ using Equation 5

Experiments

- Dataset
  - Persona-chat: 1137/99/100 users for meta-training/meta-validation/meta-testing, and each user has 121 utterances on average.
  - $D_{train}:D_{test}$ = 10:1
- Baselines
  - Seq2seq
  - Seq2seq-Finetune
  - MAML

Experiments

- Evaluation
  - Quality: Perplexity, BLEU
  - Personality: C score
    - measure the response consistency with persona description

$$\mathbf{NLI}(u, p_j) = \begin{cases} 1 & \text{if } u \text{ entails } p_j \\ 0 & \text{if } u \text{ is independent to } p_j \\ -1 & \text{if } u \text{ contradicts } p_j \end{cases}$$

$$\boldsymbol{C}(u) = \sum_{j}^{m} \mathbf{NLI}(u, p_j) \qquad (7)$$

- Results

Model	**PPL**	**BLEU**	**C**
Seq2seq	37.91	1.27	-0.16
Seq2seq-Finetune	**33.65**	1.56	-0.05
MAML	37.43	1.54	0.14
Our Model	37.31	**1.59**	**0.15**

Conclustion

- Combine meta-LSTM and MAML in personalized dialogue generation、
- Personalized adaptation of decoder
- Improve the performance of MAML
- Future work
  - More empirical studies
    - The difference of scale and learning rate among tasks
    - …

PEKING UNIVERSITY

Thanks for listening!