

ATTENTION OVER PARAMETERS FOR DIALOGUE SYSTEMS

Anonymous authors

Paper under double-blind review

ABSTRACT

Dialogue systems require a great deal of different but complementary expertise to assist, inform, and entertain humans. For example, different domains (e.g., restaurant reservation, train ticket booking) of goal-oriented dialogue systems can be viewed as different skills, and so does ordinary chatting abilities of chit-chat dialogue systems. In this paper, we propose to learn a dialogue system that independently parameterizes different dialogue skills, and learns to select and combine each of them through *Attention over Parameters (AoP)*. The experimental results show that this approach achieves competitive performance on a combined dataset of MultiWOZ (Budzianowski et al., 2018), In-Car Assistant (Eric et al., 2017), and Persona-Chat (Zhang et al., 2018). Finally, we demonstrate that each dialogue skill is effectively learned and can be combined with other skills to produce selective responses.

1 INTRODUCTION

Unlike humans who can do both, goal-oriented dialogues (Williams & Young, 2007; Young et al., 2013) and chit-chat conversations (Serban et al., 2016a; Vinyals & Le, 2015) are often learned with separate models. A more desirable approach for the users would be to have a single chat interface that can handle both casual talk and tasks such as reservation or scheduling. This can be formulated as a problem of learning different conversational skills across multiple domains. A skill can be either querying a database, generating daily conversational utterances, or interacting with users in a particular task-domain (e.g. booking a restaurant). One challenge of having multiple skills is that existing datasets either focus only on chit-chat or on goal-oriented dialogues. This is due to the fact that traditional goal-oriented systems are modularized (Williams & Young, 2007; Hori et al., 2009; Lee et al., 2009; Levin et al., 2000; Young et al., 2013); thus, they cannot be jointly trained with end-to-end architecture as in chit-chat. However, recently proposed end-to-end trainable models (Eric & Manning, 2017; Madotto et al., 2018; Wu et al., 2019; Reddy et al., 2018; Yavuz et al., 2018) and datasets (Bordes & Weston, 2017; Eric et al., 2017) allow us to combine goal-oriented (Budzianowski et al., 2018; Eric et al., 2017) and chit-chat (Zhang et al., 2018) into a single benchmark dataset with multiple conversational skills as shown in Table 1.

A straight forward solution would be to have a single model for all the conversational skills, which has shown to be effective to a certain extent by (Zhao et al., 2017) and (McCann et al., 2018). Putting aside the performance in the tasks, such fixed shared-parameter framework, without any task-specific designs, would lose controllability and interpretability in the response generation. In this paper, instead, we propose to model multiple conversational skills using the *Mixture of Experts (MoE)* (Jacobs et al., 1991) paradigm, i.e., a model that learns and combine independent specialized experts using a gating function.

For instance, each expert could specialize in different dialogues domains (e.g., Hotel, Train, Chit-Chat etc.) and skills (e.g., generate SQL query). A popular implementation of MoE (Shazeer et al., 2017; Kaiser et al., 2017) uses a set of linear transformation (i.e., experts) in between two LSTM (Schmidhuber, 1987) layers. However, several problems arise with this implementation: 1) the model is computationally expensive as it has to decode multiple times each expert and make the combination at the representation-level; 2) no prior knowledge is injected in the expert selection (e.g., domains); 3) Seq2Seq model has limited ability in extracting information from a Knowledge Base (KB) (i.e., generated by the SQL query) (Eric et al., 2017), as required in end-to-end task-oriented dialogues

Table 1: An example from the dataset which includes both chit-chat and task-oriented conversations. The model has to predict all the *Sys* turn, which includes SQL query and generating response from a the Memory content, which is dynamically updated with the queries results. The skills are the prior knowledge needed for the response, where *Persona* refers to chit-chat.

Spk.	Conversation	Skills
<i>Usr:</i>	Can you help me find a cheap 2 star hotel?	
<i>Sys:</i>	<code>SELECT * FROM hotel WHERE pricerange='cheap' AND stars=2 AND type='hotel'</code>	SQL+HOTEL
<i>Mem:</i>	(Result table from the SQL query, with real entities.)	
<i>Sys:</i>	Arms is a 2 star hotel in the centre. Shall I book a room for you?	HOTEL
<i>Usr:</i>	Yes, thank you!	
<i>Sys:</i>	Sure, for how many people should I book it for?	HOTEL
...		
<i>Sys:</i>	Hello! How are you today?	PERSONA
<i>Usr:</i>	I am good thank you! Do you like to ski? I really like it.	
<i>Sys:</i>	No, I like to code in Haskell. Do you like to code?	PERSONA

systems (Bordes & Weston, 2017). The latter can be solved by using more advanced multi-hop models like the Transformer (Vaswani et al., 2017), but the remaining two need to be addressed. Hence, in this paper we:

- propose a novel Transformer-based (Vaswani et al., 2017) architecture called *Attention over Parameters* (AoP). This model parameterize the conversational skills of end-to-end dialogue systems with independent decoder parameters (experts), and learns how to dynamically select and combine the appropriate decoder parameter sets by leveraging prior knowledge from the data such as domains and skill types;
- proof that AoP is algorithmically more efficient (Appendix A1) compared to forwarding all the Transformer decoders and then mix their output representation, like is normally done in MoE. Figure 1 illustrates the high-level intuition of the difference;
- empirically show the effectiveness of using specialized parameters in a combined dataset of Multi-WOZ (Budzianowski et al., 2018), In-Car Assistant (Eric et al., 2017), and Persona-Chat (Zhang et al., 2018), which to the best of our knowledge, is the first evaluation of this genre i.e. end-to-end large-scale multi-domains/skills. Moreover, we show that our model is highly interpretable and is able to combine different learned skills to produce compositional responses.

2 RELATED WORK

Dialogue Task-oriented dialogue models (Gao et al., 2018) can be categorized in two types: module-based (Williams & Young, 2007; Hori et al., 2009; Lee et al., 2009; Levin et al., 2000; Young et al., 2013) and end-to-end. In this paper, we focus on the latter which are systems that train a single model directly on text transcripts of dialogues. These tasks are tackled by selecting a set of predefined utterances (Bordes & Weston, 2017; Liu & Perez, 2017; Wu et al., 2017; 2018; Williams et al., 2017; Seo et al., 2017) or by generating a sequence of tokens (Madotto et al., 2018; Wen et al., 2017; Serban et al., 2016b; Zhao et al., 2017; Serban et al., 2017). Especially in the latter, copy-augmented models (Eric & Manning, 2017; Madotto et al., 2018; Wu et al., 2019; Reddy et al., 2018; Yavuz et al., 2018) are very effective since extracting entities from a knowledge base is fundamental. On the other hand, end-to-end open domain chit-chat models have been widely studied (Serban et al., 2016a; Vinyals & Le, 2015; Wolf et al., 2019). Several works improved on the initially reported baselines with various methodologies (Kulikov et al., 2018; Yavuz et al., 2018; Hancock et al., 2019; Lucas et al., 2009; Joshi et al., 2017; Zemlyanskiy & Sha, 2018; Dinan et al., 2019). Finally, (Zhao et al., 2017) was the first attempt of having an end-to-end system for both task-oriented models and chit-chat. However, the dataset used for the evaluation was small, evaluated only in single domain, and the chit-chat ability was added manually through rules.

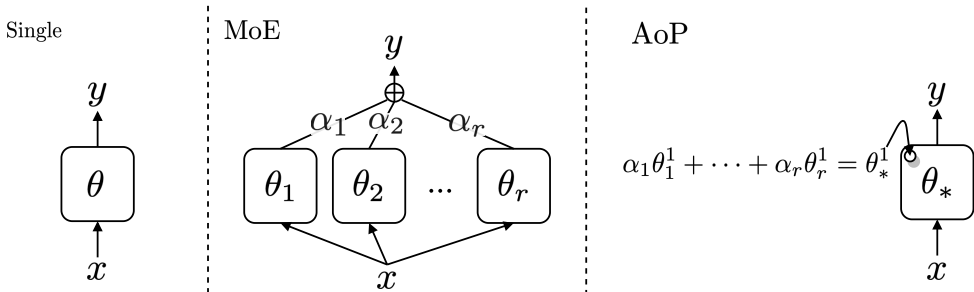


Figure 1: Comparisons between Single model, Mixture of Experts (MoE)(Jacobs et al., 1991), and Attention over Parameters (AoP).

Mixture of Expert & Conditional Computation The idea of having specialized parameters, or so-called experts, has been widely studied topics in the last two decades (Jacobs et al., 1991; Jordan & Jacobs, 1994). For instance, different architecture and methodologies have been used such as SVM (Collobert et al., 2002), Gaussian Processes (Tresp, 2001; Theis & Bethge, 2015; Deisenroth & Ng, 2015), Dirichlet Processes (Shahbaba & Neal, 2009), Hierarchical Experts (Yao et al., 2009), Infinity Number of Experts (Rasmussen & Ghahramani, 2002), and sequential expert addition (Aljundi et al., 2017). More recently, the Mixture Of Expert (Shazeer et al., 2017; Kaiser et al., 2017) model was proposed which added a large number of experts between two LSTMs (Schmidhuber, 1987). To the best of our knowledge, none of these previous works applied the results of the gating function to the parameters itself. On the other hand, there are Conditional Computational models which learn to dynamically select their computation graph (Bengio et al., 2013; Davis & Arel, 2013). Several methods have been used such as reinforcement learning (Bengio et al., 2016), a halting function (Graves, 2016; Deghani et al., 2019; Figurnov et al., 2017), by pruning (Lin et al., 2017; He et al., 2018) and routing/controller function (Rosenbaum et al., 2018). However, this line of work focuses more on optimizing the inference performance of the model more than specializing parts of it for computing a certain task.

Multi-task Learning Even though our model processes only input sequence and output sequences of text, it actually jointly learns multiple tasks (e.g. SQL and BOOK query, memory retrieval, and response generation), thus it is also related to multi-task learning (Caruana, 1997). Interested readers may refer to (Ruder, 2017; Zhou et al., 2011) for a general overview on the topic. In Natural Language Processing, multi-task learning has been applied in a wide range of applications such as parsing (Collobert et al., 2011; Hashimoto et al., 2017; Ruder et al., 2017), machine translation in multiple languages (Johnson et al., 2017), and parsing image captioning and machine translation (Luong et al., 2016). More interestingly, DecaNLP (McCann et al., 2018) has a large set of tasks that are cast to question answering (QA), and learned by a single model. In this work, we focus more on conversational data, but in future works, we plan to include these QA tasks.

3 METHODOLOGY

We use the standard encoder-decoder architecture and avoid any task-specific designs (Madotto et al., 2018; Wu et al., 2019; Reddy et al., 2018), as we aim to build a generic conversation model for both chit-chat and task-oriented dialogues. More specifically, we use a Transformer (Vaswani et al., 2017) for both encoder and decoder.

Let us define the sequence of tokens in the dialogue history as $D = \{d_1, \dots, d_m\}$ and the dynamic memory content as a sequence of tokens $M = \{m_1, \dots, m_z\}$. The latter can be the result of a SQL query execution (e.g., table) or plain texts (e.g., persona description), depending on the task. The dialogue history D and the memory M are concatenated to obtain the final input denoted by $X = [D; M] = \{x_1, \dots, x_{n=m+z}\}$. We then denote $Y = \{y_1, \dots, y_k\}$ as the sequence of tokens that the model is expected to produce. Without loss of generality, Y can be both plain text and SQL-like queries. Hence, the model has to learn when to issue database queries and when to generate human-like responses. Finally, we define a binary skill vector $V = \{v_1, \dots, v_r\}$ that specifies the

type of skills required to generate Y . This can be considered as a prior vector for learning to select the correct expert during the training¹. For example, in Table 1 the first response is of type SQL in the Hotel domain, thus the skill vector V will have $v_{\text{SQL}} = 1$ and $v_{\text{Hotel}} = 1$, while all the other skill/domains are set to zero². More importantly, we may set the vector V to have multiple ones to enforce the model to compose skills to achieve a semantic compositionality of different experts.

3.1 ENCODER-DECODER

To map the input sequence to the output sequence, we use a standard Transformer (Vaswani et al., 2017) and denote the encoder and decoder as TRS_{enc} and TRS_{dec} , respectively. The input of a Transformer is the embedded representation of the input words; thus, we define a word embedding matrix $E \in \mathbb{R}^{d \times |V|}$ where d is the embedding size and $|V|$ is the cardinality of the vocabulary. The input X , with its positional embedding (Appendix A2 for more information), are encoded as the following equation:

$$H = \text{TRS}_{\text{enc}}(E(X)), \quad (1)$$

where $H \in \mathbb{R}^{d_{\text{model}} \times n}$, and E . Then the decoder receives the target sequence shifted by one $Y_{:k-1} = \{\langle \text{SOS} \rangle, y_1, \dots, y_k\}$ as the input. Using teacher-forcing (Williams & Zipser, 1989), the model is trained to produce the correct sequence Y . The output of the decoder is produced as follow:

$$O = \text{TRS}_{\text{dec}}(E(Y_{:k-1}), H), \quad (2)$$

where $O \in \mathbb{R}^{d_{\text{model}} \times k}$. Finally, a distribution over the vocabulary is generated for each token by an affine transformation $W \in \mathbb{R}^{d_{\text{model}} \times |V|}$ followed by a Softmax function.

$$P(Y|X) = \text{Softmax}(O^T W), \quad (3)$$

In addition, $P(Y|X)$ is mixed with the encoder-decoder attention distribution to enable to copy token from the input sequence as in (See et al., 2017). The model is then trained to minimize a standard cross entropy loss function and at inference time to generate one token at the time in an auto-regressive manner (Graves, 2013). Hence, the training loss is defined as:

$$\mathcal{L}_{P(Y|X)} = - \sum_{t=1}^k \log(P(Y|X)_t(y_t)). \quad (4)$$

3.2 ATTENTION OVER PARAMETERS

The main idea is to produce a single set of parameters for decoder TRS_{dec} by the weighted sum of r independently parameterized decoders. This process is similar to attention (Luong et al., 2015) where the memories are the parameters and the query is the encoded representation. Let us define $\Theta = [\theta_1, \dots, \theta_r]$ as the list of parameters for r decoders, since a TRS_{dec} is represented by its parameters θ . Since each θ can be sized in the order of millions, we assign the corresponding key vectors to each θ , similar to key-value memory networks (Miller et al., 2016). Thus, we use a key matrix $K \in \mathbb{R}^{d_{\text{model}} \times r}$ and a Recurrent Neural Networks (RNN), in this instance a GRU (Cho et al., 2014), to produce the query vector by processing the encoder output H . The attention weights for each decoders' parameters is computed as follow:

$$q = \text{RNN}(H) \quad (5)$$

$$\alpha = \text{Softmax}(qK) \quad (6)$$

where $q \in \mathbb{R}^{d_{\text{model}}}$ and $\alpha \in \mathbb{R}^r$ is the attention vectors where each α_i is the score corresponding to θ_i . Hence, the new set of parameters is computed as follow:

$$\theta^* = \sum_i^r \alpha_i \theta_i \quad (7)$$

The combined set of parameters θ^* are then used to initialize a new TRS_{dec} , and Equation 2 will be applied to the input based on this. Equation 6 is similar to the gating function proposed

¹the vector V will be absent during the testing

²With the assumption that at each index in V is assigned a semantic skill (e.g. SQL position i)

in (Shazeer et al., 2017; Jacobs et al., 1991), but the resulting scoring vector α is applied directly to the parameter instead of the output representation of each decoder. This indeed makes the computation algorithmically faster since the forward computation is done only once and summing r elements of size $|\theta|$ is linear compared to forwarding r multiplications with the input. Interested readers may refer to Appendix A1 for the proof. Importantly, if we apply α to each of the output representation O^i generated by the TRS_{dec}^i , we end up having a Transformer-based implementation of MoE³. We call this model as Attention over Representation (AoR).

Moreover, an additional loss term is used to supervise the attention vector α by using the prior knowledge vector V . Since multiple decoder parameters can be selected at the same time, we use a binary cross-entropy to train each α_i . Thus a second loss is defined as:

$$\mathcal{L}_V = - \sum_{i=1}^r [v_i \times \log \sigma(qK)_i + (1 - v_i) \times \log(1 - \sigma(qK)_i)] \quad (8)$$

The final loss is the summation of $\mathcal{L}_{P(Y|X)}$ and \mathcal{L}_V .

Finally, in AoP, but in general in the MoE framework, stacking multiple layers (e.g., Transformer) leads to models with a large number of parameters, since multiple experts are repeated across layers. An elegant workaround is the Universal Transformer (Dehghani et al., 2019), which loops over an *unique* layer and, as shown by (Dehghani et al., 2019), holds similar or better performance than a multi-layer Transformer. In our experiment, we report a version of AoP that uses this architecture, which for instance does not add any further parameter to the model.

4 EXPERIMENTS AND RESULTS

4.1 DATASET

To evaluate the performance of our model for different conversational skills, we propose to combine three publicly available datasets: MultiWOZ (Budzianowski et al., 2018), Stanford Multi-domain Dialogue (Eric et al., 2017) and Persona-Chat (Zhang et al., 2018).

MultiWOZ (MWOZ) is a human-to-human multi-domain goal-oriented dataset annotated with dialogue acts and states. In this dataset, there are seven domains (i.e., Taxi, Police, Restaurant, Hospital, Hotel, Attraction, Train) and two APIs interfaces: SQL and BOOK. The former is used to retrieve information about a certain domain and the latter is used to book

restaurants, hotels, trains, and taxis. We refine this dataset to include SQL/BOOK queries and their outputs using the same annotations schema as (Bordes & Weston, 2017). Hence, each response can either be plain text conversation with the user or SQL/BOOK queries, and the memory is dynamically populated with the results from the queries as the generated response is based on such information. This transformation allows us to train end-to-end models that learns how and when to produce SQL queries, to retrieve knowledge from a dynamic memory, and to produce plain text response. A detailed explanation is reported in Appendix A4, together with some samples.

Stanford Multi-domain Dialogue (SMD) is another human-to-human multi-domain goal-oriented dataset that is already designed for end-to-end training. There are three domains in this dataset (i.e., Point-of-Interest, Weather, Calendar). The difference between this dataset and MWOZ is that each dialogue is associated with a set of records relevant to the dialogues. The memory is fixed in this case so the model does not need to issue any API calls. However, retrieving the correct entities from the memory is more challenging as the model has to compare different alternatives among records.

Persona-Chat is a multi-turn conversational dataset, in which two speakers are paired and different persona descriptions (4-5 sentences) are randomly assigned to each of them. For example, “*I am*

³but with the additional prior over the expert

Figure 2: Datasets statistics

	SMD	MWOZ	Persona
#Dialogues	2425	8,438	12,875
#turns	12,732	115,424	192,690
Avg. turns	5.25	13.68	14.97
Avg. tokens	8.02	13.18	11.96
Vocab	2,842	24,071	20,343

Table 2: Results for the goal-oriented responses in both MWOZ and SMD. Last row, and italicized, are the Oracle results, and bold-faced are best in each setting (w and w/o Universal). Results are averaged among three run (full table in Appendix A6).

Model	F1	BLEU	SQL _{Acc}	SQL _{BLEU}	BOOK _{Acc}	BOOK _{BLEU}
<i>Seq2Seq</i>	38.37	9.42	49.97	81.75	39.05	79.00
<i>TRS</i>	36.91	9.92	61.96	89.08	46.51	78.41
<i>MoE</i>	38.64	9.47	53.60	85.38	37.23	78.55
<i>AoR</i>	40.36	10.66	69.39	90.64	52.15	81.15
<i>AoP</i>	42.26	11.14	71.1	90.90	56.31	84.08
<i>TRS + U</i>	39.39	9.29	61.80	89.70	50.16	79.05
<i>AoP + U</i>	44.04	11.26	74.83	91.90	56.37	84.15
<i>AoP w/o \mathcal{L}_V</i>	38.50	10.50	61.47	88.28	52.61	80.34
<i>AoP+O</i>	46.36	11.99	73.41	93.81	56.18	86.42

an old man” and *“I like to play football”* are one of the possible persona descriptions provided to the system. Training models using this dataset results in a more persona consistent and fluent conversation compared to other existing datasets (Zhang et al., 2018). Currently, this dataset has become one of the standard benchmarks for chit-chat systems, thus, we include it in our evaluation.

For all three datasets, we use the training/validation/test split provided by the authors and we keep all the *real entities* in input instead of using their delexicalized version as in (Budzianowski et al., 2018; Eric et al., 2017). This makes the task more challenging, but at the same time more interesting since we force the model to produce real entities instead of generic and frequent placeholders. Table 2 summarizes the dataset statistics in terms of number of dialogues, turns, and unique tokens. Finally, we merge the three datasets obtaining 154,768/19,713/19,528 for training, validation and, test respectively, and a vocabulary size of 37,069 unique tokens.

4.2 EVALUATION METRICS

Goal-Oriented For both MWOZ and SMD, we follow the evaluation done by existing works (Eric & Manning, 2017; Zhao et al., 2017; Madotto et al., 2018; Wu et al., 2017). We use BLEU⁴ score (Papineni et al., 2002) to measure the response fluency and Entity F1-Score (Wen et al., 2017; Zhao et al., 2017) to evaluates the ability of the model to generate relevant entities from the dynamic memory. Since MWOZ also includes SQL and BOOK queries, we compute the exact match accuracy (i.e., ACC_{SQL} and ACC_{BOOK}) and BLEU score (i.e., $BLEU_{SQL}$ and $BLEU_{BOOK}$). Furthermore, we also report the F1-score for each domain in both MWOZ and SMD.

Chit-Chat We compare perplexity, BLEU score, F1-score (Dinan et al., 2019), and Consistency score of the generate sentences with the human-generated prediction. The Consistency score is computed using a Natural Language Inference (NLI) model trained on dialogue NLI (Sean et al., 2018), a recently proposed corpus based on Persona dataset. We fine-tune a pre-trained BERT model (Devlin et al., 2018) using the dialogue DNLi corpus and achieve a test set accuracy of 88.43%, which is similar to the best-reported model in (Sean et al., 2018). The consistency score is defined as follow:

$$\text{NLI}(u, p_j) = \begin{cases} 1 & \text{if } u \text{ entails } p_j \\ 0 & \text{if } u \text{ is independent of } p_j \\ -1 & \text{if } u \text{ contradicts } p_j \end{cases} \quad C(u) = \sum_j^m \text{NLI}(u, p_j) \quad (9)$$

where u is a generated utterance and p_j is one sentence in the persona description. In Sean et al. (2018), the authors showed that by re-ranking the beam search hypothesis using the DNLi score (i.e., C score), they achieved a substantial improvement in dialogue consistency. Intuitively, having a higher consistency C score means having a more persona consistent dialogue response.

⁴Using the `multi-bleu.perl` script

4.3 BASELINES

In our experiments, we compare Sequence-to-Sequence (*Seq2Seq*) (See et al., 2017), Transformer (*TRS*) (Vaswani et al., 2017), Mixture of Expert (*MoE*) (Shazeer et al., 2017) and Attention over Representation (*AoR*) with our proposed Attention over Parameters (*AoP*). In all the models, we used the same copy-mechanism as in (See et al., 2017). In *AoR* instead of mixing the parameters as in Equation 7, we mix the output representation of each transformer decoder (i.e. Equation 2). For all *AoP*, *AoR*, and *MoE*, $r = 13$ is the number of decoders (experts): 2 skills of SQL and BOOK, 10 different domains for MWOZ+SMD, and 1 for Persona-Chat. Furthermore, we include also the following experiments: *AoP* that uses the gold attention vector V , which we refer as *AoP w/ Oracle* (or *AoP + O*); *AoP* trained by removing the \mathcal{L}_V from the optimization (*AoP w/o \mathcal{L}_V*); and as aforementioned, the Universal Transformer for both *AoP* (*AoP + U*) and the standard Transformer (*TRS + U*) (i.e., 6 hops). All detailed model description and the full set of hyper-parameters used in the experiments are reported in Appendix A5.

4.4 RESULTS

Table 2 and Table 3 show the respectively evaluation results in MWOZ+SMD and Persona-Chat datasets. From Table 2, we can identify four patterns. 1) *AoP* and *AoR* perform consistently better than other baselines which shows the effectiveness of combining parameters by using the correct prior V ; 2) *AoP* performs consistently, but marginally, better than *AoR*, with the advantage of an algorithmic faster inference; 3) Using Oracle (*AoP+O*) gives the highest performance in all the measures, which shows the performance upper-bound for *AoP*. Hence, the performance gap when not using oracle attention is most likely due to the error in attention α (i.e., 2% error rate). Moreover, Table 2 shows that by removing \mathcal{L}_V (*AoP w/o \mathcal{L}_V*) the model performance decreases, which confirms that good inductive bias is important for learning how to select and combine different parameters (experts). Additionally, in Appendix A6, we report the per-domain F1-Score for SQL, BOOK and sentences, and Table 3 and Table 2 with the standard deviation among the three runs.

Furthermore, from Table 3, we can notice that *MoE* has the lowest perplexity and F1-score, but *AoP* has the highest Consistency and BLEU score. Notice that the perplexity reported in (Zhang et al., 2018) is lower since the vocabulary used in their experiments is smaller. In general, the difference in performance among different models is marginal except for the Consistency score; thus, we can conclude that all the models can learn this skill reasonably well. Consistently with the previous results, when \mathcal{L}_V is removed from the optimization, the models' performance decreases.

Finally, in both Table 2 and Table 3, we report the results obtained by using the Universal Transformer, for both *AoP* and the Transformer. By adding the layer recursion, both models are able to consistently improve all the evaluated measures, in both Persona-Chat and the Task-Oriented tasks. Especially *AoP*, which achieves better performance than Oracle (i.e. single layer) in the SQL accuracy, and a consistently better performance in the Persona-Chat evaluation.

5 SKILL COMPOSITION

To demonstrate the effectiveness of our model in learning independent skills and composing them together, we *manually trigger* skills by modifying α and generate 14 different responses for the same input dialogue context. This experiment allows us to verify whether the model accurately captures the meaning of each skill and whether it can properly learn to compose the selected parameters (skills). Table 3 first shows the dialogue history along with the response of *AoP* on the top, and then different responses generated by modifying α (i.e., black cells correspond to 1 in the vector, while the whites are 0). By analyzing Table 3⁵ we can notice that:

Figure 3: Results for the Persona-Chat dataset.

Model	Ppl.	F1	C	BLEU
<i>Seq2Seq</i>	39.42	6.33	0.11	2.79
<i>TRS</i>	43.12	7.00	0.07	2.56
<i>MoE</i>	38.63	7.33	0.19	2.92
<i>AoR</i>	40.18	6.66	0.12	2.69
<i>AoP</i>	39.14	7.00	0.21	3.06
<i>TRS + U</i>	43.04	7.33	0.15	2.66
<i>AoP + U</i>	37.40	7.00	0.29	3.22
<i>AoP w/o \mathcal{L}_V</i>	42.81	6.66	0.12	2.85
<i>AoP + O</i>	<i>40.16</i>	<i>7.33</i>	<i>0.21</i>	<i>2.91</i>

Table 3: Selecting different skills thought the attention vector α results in a skill-consistent response. *AoP* response activates SQL and Train.

Sys: There are lots of trains to choose from! Where are you departing from?

Usr: I am departing from london heading to cambridge.

Sys: What time will you be travelling?

Usr: I need to arrive by 1530.

AoP: SELECT * FROM train WHERE destination="cambridge" AND day="monday" AND arriveBy < "1530" and departure="london"

SQL	BOOK	Taxi	Restaurant	Hotel	Attraction	Train	Schedule	Persona	<i>AoP</i> responses using different parameter combination
		■							I would be happy to book that for you.
			■						I have found you a few to choose from. what time would ..
				■					I have many options for you. is there a certain time you ...
					■				Is there a certain area you would like?
						■			what time would you like to leave?
							■		okay I have two trains what time would you like to cambridge?
								■	Where do you do for work ?
	■	■							BOOK FROM taxi WHERE leaveAt>"1530" AND destination=..
	■		■						BOOK FROM restaurant WHERE time="1530" AND ...
	■			■					BOOK FROM hotel WHERE people="1" AND day="monday"
	■					■			BOOK FROM train WHERE people="1" AND id_booking='
■			■						SELECT * FROM restaurant WHERE food="1530" ...
■				■					SELECT * FROM hotel WHERE type="1530"
■					■				SELECT * FROM attraction WHERE name="departure"

- The model learns the correct semantics of each skill. For instance, the *AoP* response is of type SQL and Train, and by deactivating the SQL skill and activating other domain-skills, including Train, we can see that the responses are grammatical and they are coherent with the selected skill semantics. For instance, by just selecting *Train*, the generated answer becomes “*what time would you like to leave?*” which is coherent with the dialogue context since such information has not been yet provided. Interestingly, when *Persona* skill is selected, the generated response is conversational and also coherent with the dialogue, even though it is less fluent.
- The model effectively learns how to compose multiple skills. For instance, when SQL or BOOK are triggered the response produces the correct SQL-syntax (e.g. “SELECT * FROM ..” etc.). By also adding the corresponding domain-skill, the model generates the correct query format and attributes relative to the domain type (e.g. in *SQL*, *Restaurant*, the model queries with the relevant attribute *food* for restaurants).

6 CONCLUSION

In this paper, we propose a novel way to train a single end-to-end dialogue model with multiple composable and interpretable skills. Unlike previous work, that mostly focused on the representation-level mixing (Shazeer et al., 2017), our proposed approach, *Attention over Parameters*, learns how to softly combine independent sets of specialized parameters (i.e., making SQL-Query, conversing with consistent persona, etc.) into a single set of parameters. By doing so, we not only achieve compositionality and interpretability but also gain algorithmically faster inference speed. To train and evaluate our model, we organize a multi-domain task-oriented datasets into end-to-end trainable formats and combine it with a conversational dataset (i.e. Persona-Chat). Our model learns to consider each task and domain as a separate skill that can be composed with each other, or used independently, and we verify the effectiveness of the interpretability and compositionality with competitive experimental results and thorough analysis.

REFERENCES

- Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3366–3375, 2017.
- Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. *ICLR*, 2016.
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Antoine Bordes and Jason Weston. Learning end-to-end goal-oriented dialog. *International Conference on Learning Representations*, abs/1605.07683, 2017.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 5016–5026, 2018.
- Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, 2014.
- Ronan Collobert, Samy Bengio, and Yoshua Bengio. A parallel mixture of svms for very large scale problems. In *Advances in Neural Information Processing Systems*, pp. 633–640, 2002.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12 (Aug):2493–2537, 2011.
- Andrew Davis and Itamar Arel. Low-rank approximations for conditional feedforward computation in deep neural networks. *arXiv*, 2013.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. *ICLR*, 2019.
- Marc Deisenroth and Jun Wei Ng. Distributed gaussian processes. In *International Conference on Machine Learning*, pp. 1481–1490, 2015.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. The second conversational intelligence challenge (convai2). *arXiv preprint arXiv:1902.00098*, 2019.
- Mihail Eric and Christopher Manning. A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 468–473, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/E17-2075>.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 37–49. Association for Computational Linguistics, 2017. URL <http://aclweb.org/anthology/W17-5506>.
- Michael Figurnov, Maxwell D Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1039–1048, 2017.

- Jianfeng Gao, Michel Galley, and Lihong Li. Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 1371–1374. ACM, 2018.
- Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. Learning from dialogue after deployment: Feed yourself, chatbot! *arXiv preprint arXiv:1901.05415*, 2019.
- Kazuma Hashimoto, Yoshimasa Tsuruoka, Richard Socher, et al. A joint many-task model: Growing a neural network for multiple nlp tasks. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1923–1933, 2017.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 784–800, 2018.
- Chiori Hori, Kiyonori Ohtake, Teruhisa Misu, Hideki Kashioka, and Satoshi Nakamura. Statistical dialog management applied to wfst-based dialog systems. In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2009. ICASSP 2009.*, pp. 4793–4796. IEEE, 2009.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, Geoffrey E Hinton, et al. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351, 2017.
- Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.
- Chaitanya K Joshi, Fei Mi, and Boi Faltings. Personalization in goal-oriented dialog. *arXiv preprint arXiv:1706.07503*, 2017.
- Lukasz Kaiser, Aidan N Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. One model to learn them all. *arXiv*, 2017.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Ilya Kulikov, Alexander H Miller, Kyunghyun Cho, and Jason Weston. Importance of a search strategy in neural dialogue modelling. *arXiv preprint arXiv:1811.00907*, 2018.
- Cheongjae Lee, Sangkeun Jung, Seokhwan Kim, and Gary Geunbae Lee. Example-based dialog modeling for practical multi-domain dialog system. *Speech Communication*, 51(5):466–484, 2009.
- Esther Levin, Roberto Pieraccini, and Wieland Eckert. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on speech and audio processing*, 8(1): 11–23, 2000.
- Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *Advances in Neural Information Processing Systems*, pp. 2181–2191, 2017.
- Fei Liu and Julien Perez. Gated end-to-end memory networks. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pp. 1–10, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/E17-1001>.
- JM Lucas, F Fernández, J Salazar, J Ferreiros, and R San Segundo. Managing speaker identity and user profiles in a spoken dialogue system. *Procesamiento del lenguaje natural*, 1(43), 2009.

- Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1412–1421, Lisbon, Portugal, September 2015. Association for Computational Linguistics. URL <http://aclweb.org/anthology/D15-1166>.
- Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. Multi-task sequence to sequence learning. In *International Conference on Learning Representations*, 2016.
- Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1468–1478, 2018.
- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. The natural language decathlon: Multitask learning as question answering. *arXiv preprint arXiv:1806.08730*, 2018.
- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1400–1409, 2016.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL <http://www.aclweb.org/anthology/P02-1040>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Carl E Rasmussen and Zoubin Ghahramani. Infinite mixtures of gaussian process experts. In *Advances in neural information processing systems*, pp. 881–888, 2002.
- Revanth Reddy, Danish Contractor, Dinesh Raghu, and Sachindra Joshi. Multi-level memory for task oriented dialogs. *arXiv preprint arXiv:1810.10647*, 2018.
- Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. Routing networks: Adaptive selection of non-linear functions for multi-task learning. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=ry8dvM-R->.
- Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.
- Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. Learning what to share between loosely related tasks. *arXiv preprint arXiv:1705.08142*, 2017.
- Jurgen Schmidhuber. Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook. Diploma thesis, Technische Universitat Munchen, Germany, 14 May 1987. URL <http://www.idsia.ch/~juergen/diploma.html>.
- Welleck Sean, Jason Weston, Arthur Szlam, and Kyunghyun Cho. Dialogue natural language inference. *arXiv preprint arXiv:1811.00671*, 2018.
- Abigail See, Peter J. Liu, and Christopher D. Manning. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1073–1083, Vancouver, Canada, July 2017. Association for Computational Linguistics. URL <http://aclweb.org/anthology/P17-1099>.
- Minjoon Seo, Sewon Min, Ali Farhadi, and Hannaneh Hajishirzi. Query-reduction networks for question answering. *International Conference on Learning Representations*, 2017.
- Iulian Vlad Serban, Ryan Lowe, Laurent Charlin, and Joelle Pineau. Generative deep neural networks for dialogue: A short review. *arXiv preprint arXiv:1611.06216*, 2016a.

- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, pp. 3776–3784, 2016b.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C Courville, and Yoshua Bengio. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*, pp. 3295–3301, 2017.
- Babak Shahbaba and Radford Neal. Nonlinear models using dirichlet process mixtures. *Journal of Machine Learning Research*, 10(Aug):1829–1850, 2009.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- Lucas Theis and Matthias Bethge. Generative image modeling using spatial lstms. In *Advances in Neural Information Processing Systems*, pp. 1927–1935, 2015.
- Volker Tresp. Mixtures of gaussian processes. In *Advances in neural information processing systems*, pp. 654–660, 2001.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 6000–6010, 2017.
- Oriol Vinyals and Quoc V Le. A neural conversational model. *arXiv preprint arXiv:1506.05869*, 2015.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Lina Maria Rojas-Barahona, Pei hao Su, Stefan Ultes, David Vandyke, and Steve J. Young. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*, 2017.
- Jason D Williams and Steve Young. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422, 2007.
- Jason D Williams, Kavosh Asadi, and Geoffrey Zweig. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 665–677, Vancouver, Canada, July 2017. Association for Computational Linguistics. URL <http://aclweb.org/anthology/P17-1062>.
- Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*, 2019.
- Chien-Sheng Wu, Andrea Madotto, Genta Winata, and Pascale Fung. End-to-end recurrent entity network for entity-value independent goal-oriented dialog learning. In *Dialog System Technology Challenges Workshop, DSTC6*, 2017.
- Chien-Sheng Wu, Andrea Madotto, Genta Winata, and Pascale Fung. End-to-end dynamic query memory network for entity-value independent task-oriented dialog. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- Chien-Sheng Wu, Richard Socher, and Caiming Xiong. Global-to-local memory pointer networks for task-oriented dialogue. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Bangpeng Yao, Dirk Walther, Diane Beck, and Li Fei-Fei. Hierarchical mixture of classification experts uncovers interactions between brain regions. In *Advances in Neural Information Processing Systems*, pp. 2178–2186, 2009.

- Semih Yavuz, Abhinav Rastogi, Guan-lin Chao, Dilek Hakkani-Tür, and Amazon Alexa AI. Deep-copy: Grounded response generation with hierarchical pointer networks. *Conversational AI NIPS workshop*, 2018.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179, 2013.
- Yury Zemlyanskiy and Fei Sha. Aiming to know you better perhaps makes me a more engaging dialogue partner. *CoNLL 2018*, pp. 551, 2018.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2204–2213. Association for Computational Linguistics, 2018. URL <http://aclweb.org/anthology/P18-1205>.
- Tiancheng Zhao, Allen Lu, Kyusong Lee, and Maxine Eskenazi. Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 27–36. Association for Computational Linguistics, August 2017. URL <http://aclweb.org/anthology/W17-5505>.
- Jiayu Zhou, Jianhui Chen, and Jieping Ye. Clustered multi-task learning via alternating structure optimization. In *Advances in neural information processing systems*, pp. 702–710, 2011.

A APPENDIX

A.1 COMPUTATIONAL COST AOP

Corollary A.0.1. *The computation cost of Attention over Parameters (AoP) is always lower than Mixture Of Experts (MoE) as long as the processed sequence is longer than 1.*

Proof. Let $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^n$ a generic function parametrized by θ . Without loss of generality, we define θ as a affine transformation $W \in \mathbb{R}^{d \times n}$. Let $X \in \mathbb{R}^{t \times d}$ a generic input sequence of length t and d dimensional size. Let the set $F = [f_{\theta_1}, \dots, f_{\theta_r}]$ be the set of r experts. Hence, the operation done by MoE are:

$$\text{MoE}(X) = f_{\theta_1}(X) + \dots + f_{\theta_r}(X) = XW_1 + \dots + XW_r \tag{10}$$

Thus the computational cost in term of operation is $\mathcal{O}(rtdn + rtn)$ since the cost of $f_{\theta_i}(X)$ is $\mathcal{O}(tdn)$ and it is repeated r times, and the cost of summing the representation is $\mathcal{O}(rtn)$. On the other hand, the operation done by AoP are:

$$\theta^* = \theta_1 + \dots + \theta_r = W_1 + \dots + W_r \tag{11}$$

$$\text{AoP}(X) = f_{\theta^*}(X) = XW^* \tag{12}$$

in this case the computational cost in term of operation is $\mathcal{O}((r + t)dn)$ since the cost of summing the parameters is $\mathcal{O}(rdn)$ and the cost of f_{θ^*} is $\mathcal{O}(tdn)$. Hence, it is easy to verify that if $t > 1$ then:

$$rtdn + rtn \geq (rt)dn \geq (r + t)dn \tag{13}$$

Furthermore, the assumption of using a simple affine transformation W is actually an optimal case. Indeed, assuming that the cost of parameters sum is equal to the number of operation is optimistic, for instance already by using attention the number of operations increases but the number of parameters remains constant. \square

A.2 EMBEDDED REPRESENTATION

Since the model input may include structured data (e.g. DB records) we further define another embedding matrix for encoding the types and the segments as $P \in \mathbb{R}^{d \times |S|}$ where S is the set of positional tokens and $|S|$ its cardinality. P is used to inform the model of the token types such as speaker information (e.g. *Sys* and *Usr*), the data-type for the memory content (e.g. *Miles*, *Traffic* etc.), and segment types like dialogue turn information and database record index (Wolf et al., 2019). Figure 4 shows an example of the embedded representation of the input. Hence, we denote X_T and X_R as the type and segment tokens for each token in input X , respectively.

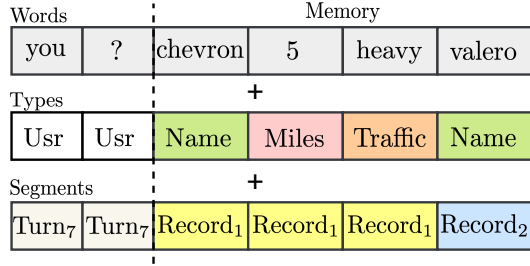


Figure 4: Positional Embedding of the dialogue history and the memory content.

A.3 ATTENTION VISUALIZATION

Figure 5 shows the attention vector α over parameters for different generated sentences. In this figure, and by analyzing more examples ⁵, we can identify two patterns:

- *AoP* learns to focus on the correct skills (i.e., SQL, BOOK) when API-calls are needed. From the first example in Figure 5, we can see that the activations in α are consistent with those in the correct attention vector P . There are also false positives, in which *AoP* puts too high weights on BOOK when the correct response is plain text that should request more information from the user (i.e., *i can help you with that. when would you like to leave the hotel?*). However, we can notice that this example is, in fact, "almost correct" as triggering a booking API call may also be considered a valid response. Meanwhile, the third example also fails to attend to the correct skill,

⁵Available in supplementary material and later online.

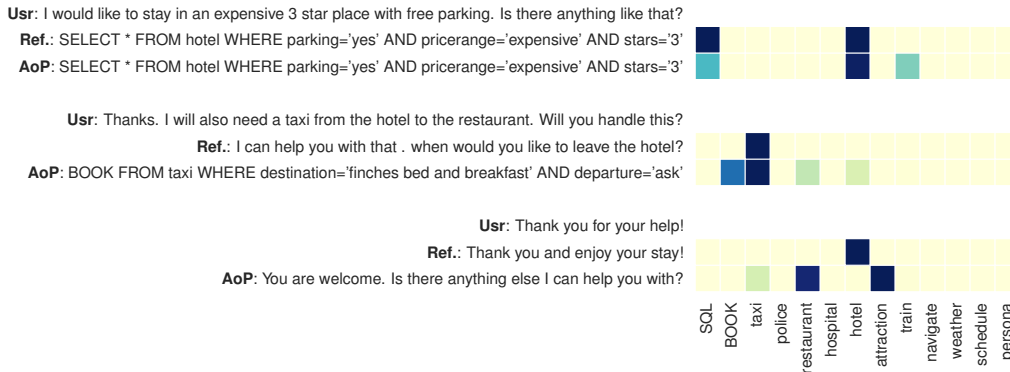


Figure 5: Attention over Parameters visualization, vector α for different reference (**Ref.**) and *AoP* generated answers. Top rows (**Usr**) are the last utterances from each dialogue contexts.

but, in fact, generates a very fluent and relevant response. This is most likely because the answer is simple and generic.

- The attention often focuses on multiple skills not directly relevant to the task. We observe this pattern especially when there are other skill-related entities mentioned in the context or the response. For example, in the second dialog example in Figure 5, we can notice that *AoP* not only accurately focuses on *taxi* domain, but also has non-negligible activations for *restaurant* and *hotel*. This is because the words “hotel” and “restaurant” are both mentioned in the dialogue context and the model has to produce two entities of the same type (i.e. *finches bed and breakfast* and *ask*).

A.4 DATA PRE-PROCESSING

As mentioned in the main article, we convert MultiWOZ into an end-to-end trainable dataset. This requires to add sql-syntax queries when the system includes particular entities. To do so we leverage two annotations such as the state-tracker and the speech acts. The first is used to generate the a well-formed query, including key and attribute, the second instead to decide when to include the query. More details on the dialogue state-tracker slots and slots value, and the different speech acts can be found in (Budzianowski et al., 2018).

A query is create by the slots, and its values, that has been updated in the latest turn. The SQL query uses the following syntax:

SELECT * FROM domain WHERE [slot_type = slot_value]*

Similarly for the booking api BOOK the syntax is the following:

BOOK FROM domain WHERE [slot_type = slot_value]*

In both cases the slot values are kept as real entities.

More challenging is to decide when to issue such apis. Speech acts are used to decide by using the "INFORM-DOMAIN" and "RECOMMEND-DOMAIN" tag. Thus any response that include those speech tag will trigger an api if and only if:

- there has been a change in the state-tracker from the previous turn
- the produced query has never been issued before

By a manual checking, this strategy results to be effective. However, as reported by (Budzianowski et al., 2018) the speech act annotation includes some noise, which is reflected also into our dataset.

The results from the SQL query can be of more that 1K records with multiple attributes. Following (Budzianowski et al., 2018) we use the following strategy:

- If no speech act INFORM or RECOMMEND and the number of records are more than 5, we use a special token in the memory $\langle TM \rangle$.

- If no speech act INFORM or RECOMMEND and the number of records are less or equal than 5, we put all the records in memory.
- If any speech act INFORM or RECOMMEND, we filter the records to include based on the act value. Notice that this is a fair strategy, since all the resulting record are correct possible answers and the annotators pick-up on of the record randomly (Budzianowski et al., 2018).

Notice that the answer of a booking call instead, is only one record containing the booking information (e.g. reference number, taxi plate etc.) or "Not Available" token in case the booking cannot made.

A.5 HYPER-PARAMETERS AND TRAINING

We used a standard Transformer architecture (Vaswani et al., 2017) with pre-trained GloVe embedding (Pennington et al., 2014). For the both Seq2Seq and MoE we use Adam (Kingma & Ba, 2014) optimizer with a learning rate of 1×10^{-3} , where instead for the Transformer we used a warm-up learning rate strategy as in (Vaswani et al., 2017). In both *AoP* and *AoR* we use an additional transformer layer on top the output of the model. Figure 6,7,8 shows the high level design MoE, AoR and AoP respectively. In all the model we used a batch size of 16, and we early stopped the model using the Validation set. All the experiments has been conducted using a single Nvidia 1080ti.

We used a small grid-search for tuning each model. The selected hyper-parameters are reported in Table 4, and we run each experiment 3 times and report the mean and standard deviation of each result.

<i>Model</i>	<i>d</i>	<i>d_{model}</i>	<i>Layers</i>	<i>Head</i>	<i>Depth</i>	<i>Filter</i>	<i>GloVe</i>	<i>Experts</i>
<i>Seq2Seq</i>	100	100	1	-	-	-	Yes	-
<i>TRS</i>	300	300	1	2	40	50	Yes	-
<i>MoE</i>	100	100	2	-	-	-	Yes	13
<i>AoP,AoR</i>	300	300	1	2	40	50	Yes	13
<i>TRS/AoP+U</i>	300	300	6	2	40	50	Yes	13

Table 4: Hyper-Parameters used for the evaluations.

A.6 MWOZ AND SMD WITH STD.

Model	F1	BLEU	SQL _{Acc}	SQL _{BLEU}	BOOK _{Acc}	BOOK _{BLEU}
<i>Seq2Seq</i>	38.37 ± 1.69	9.42 ± 0.38	49.97 ± 3.49	81.75 ± 2.54	39.05 ± 9.52	79.00 ± 3.63
<i>TRS</i>	36.91 ± 1.24	9.92 ± 0.43	61.96 ± 3.95	89.08 ± 1.23	46.51 ± 5.46	78.41 ± 2.03
<i>MoE</i>	38.64 ± 1.11	9.47 ± 0.59	53.60 ± 4.62	85.38 ± 2.68	37.23 ± 3.89	78.55 ± 2.62
<i>AoR</i>	40.36 ± 1.39	10.66 ± 0.34	69.39 ± 1.05	90.64 ± 0.83	52.15 ± 2.22	81.15 ± 0.32
<i>AoP</i>	42.26 ± 2.39	11.14 ± 0.39	71.1 ± 0.47	90.90 ± 0.81	56.31 ± 0.46	84.08 ± 0.99
<i>TRS + U</i>	39.39 ± 1.23	9.29 ± 0.71	61.80 ± 4.82	89.70 ± 1.40	50.16 ± 1.18	79.05 ± 1.42
<i>AoP + U</i>	44.04 ± 0.92	11.26 ± 0.07	74.83 ± 0.79	91.90 ± 1.03	56.37 ± 0.92	84.15 ± 0.32
<i>AoP w/o L_V</i>	38.50 ± 1.15	10.50 ± 0.55	61.47 ± 0.15	88.28 ± 0.50	52.61 ± 0.56	80.34 ± 0.21
<i>AoP+O</i>	46.36 ± 0.92	11.99 ± 0.03	73.41 ± 0.59	93.81 ± 0.16	56.18 ± 1.55	86.42 ± 0.92

A.7 PERSONA RESULT WITH STD

Model	Ppl.	F1	C	BLEU
<i>Seq2Seq</i>	39.42 ± 1.54	6.33 ± 0.58	0.11 ± 0.06	2.80 ± 0.09
<i>TRS</i>	43.12 ± 1.46	7.00 ± 0.00	0.07 ± 0.16	2.56 ± 0.07
<i>MoE</i>	38.63 ± 0.20	7.33 ± 0.05	0.19 ± 0.16	2.92 ± 0.48
<i>AoR</i>	40.18 ± 0.74	6.66 ± 0.05	0.12 ± 0.14	2.69 ± 0.34
<i>AoP</i>	39.14 ± 0.48	7.00 ± 0.00	0.21 ± 0.05	3.06 ± 0.08
<i>TRS + U</i>	43.04 ± 1.78	7.33 ± 0.57	0.15 ± 0.02	2.66 ± 0.43
<i>AoP + U</i>	37.40 ± 0.08	7.00 ± 0.00	0.29 ± 0.07	3.22 ± 0.04
<i>AoP w/o L_V</i>	42.81 ± 0.01	6.66 ± 0.57	0.12 ± 0.04	2.85 ± 0.21
<i>AoP + O</i>	40.16 ± 0.56	7.33 ± 0.58	0.21 ± 0.14	2.98 ± 0.05

A.8 DOMAIN F1-SCORE

Sentence	Seq2Seq	MoE	TRS	AoR	AoP	Aop+O
<i>Taxi</i>	71.77	75.97	73.92	76.07	76.58	78.30
<i>Police</i>	49.73	49.95	50.24	51.95	56.61	52.05
<i>Restaurant</i>	50.20	49.59	48.34	50.58	50.47	50.90
<i>Hotel</i>	46.82	45.37	43.38	45.51	46.40	44.47
<i>Attraction</i>	37.87	35.21	33.10	36.97	38.79	37.51
<i>Train</i>	46.02	41.72	41.28	44.33	46.32	45.93
<i>Weather</i>	40.38	27.06	18.97	44.77	51.94	55.23
<i>Schedule</i>	35.98	43.94	38.95	32.90	54.18	52.99
<i>Navigate</i>	18.57	21.34	6.96	12.69	12.18	16.56
BOOK						
<i>Taxi</i>	23.16	32.28	30.70	41.93	46.66	43.15
<i>Restaurant</i>	45.02	28.26	49.72	55.70	58.51	57.70
<i>Hotel</i>	49.22	31.48	41.61	51.46	56.62	57.41
<i>Train</i>	55.86	56.38	57.51	57.51	59.15	60.80
SQL						
<i>Police</i>	81.33	0.00	90.66	76.00	93.33	100.0
<i>Restaurant</i>	71.58	68.00	75.90	81.27	80.43	84.15
<i>Hospital</i>	62.22	15.55	58.89	71.11	76.67	83.33
<i>Hotel</i>	45.25	42.09	48.61	56.69	59.75	63.75
<i>Attraction</i>	65.48	67.69	65.91	70.61	76.22	74.93
<i>Train</i>	30.02	41.01	55.67	66.61	67.34	69.50

Table 5: Per Domain F1 Score.

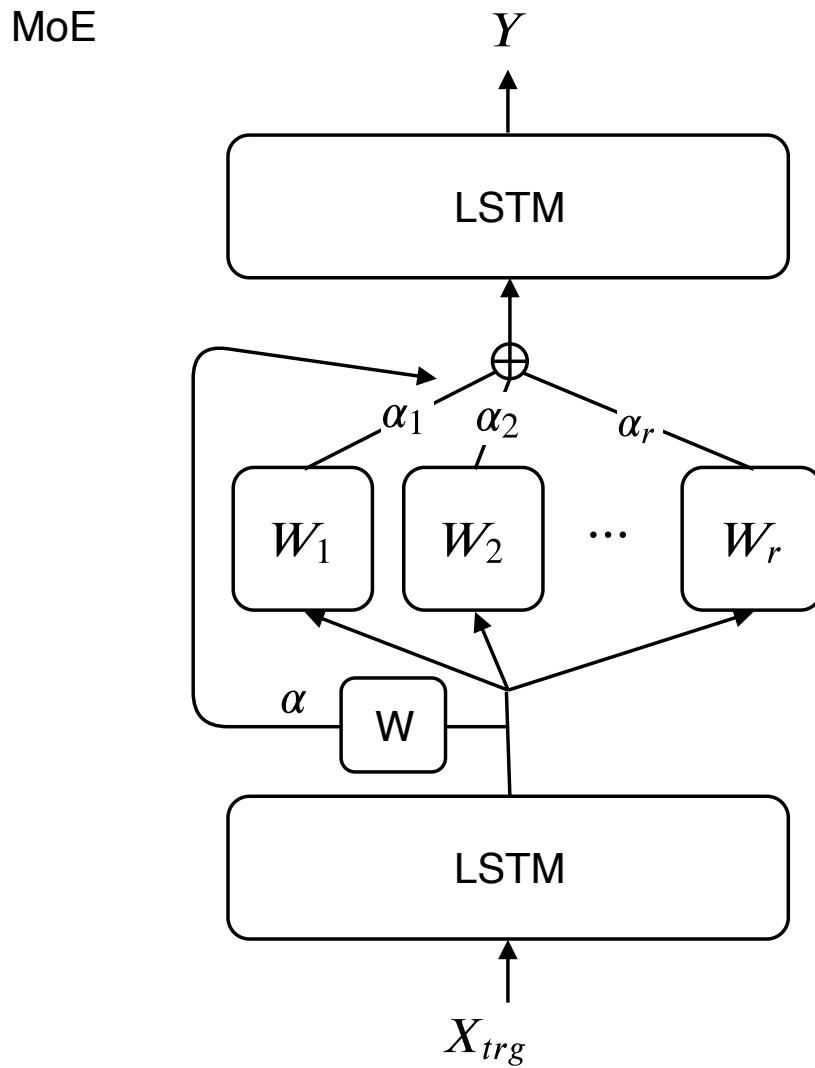


Figure 6: Mixture of Experts (MoE) (Shazeer et al., 2017) model consist of r feed-forward neural network (experts) which are embedded between two LSTM layers, a trainable gating network to select experts.

AoR




Figure 7: Attention over Representation (AoR) consist of a transformer encoder which encode the source input and compute the attention over the skills. Then r transformer decoder layers computes r specialized representation and the output response is generated based on the weighted sum the representation. In the figure, we omitted the output layer.

AoP

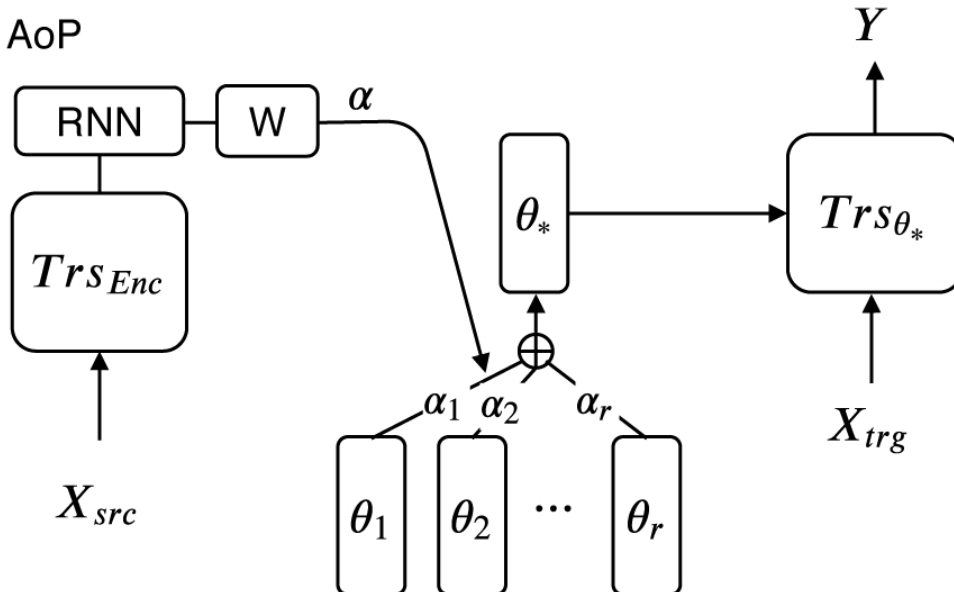


Figure 8: Attention over Parameters (AoP) consist of a transformer encoder which encode the source input and compute the attention over the skills. Then, r specialized transformer decoder layers and a dummy transformer decoder layer parameterized by the weighted sum of the r specialized transformer decoder layers parameters. In the figure, we omitted the output layer.