

Task-Oriented Dialog Systems that Consider Multiple Appropriate Responses under the Same Context

Yichi Zhang*
Tsinghua University
Beijing, China
zhangyic17@mails.tsinghua.edu.cn

Zhijian Ou
Tsinghua University
Beijing, China
ozj@tsinghua.edu.cn

Zhou Yu
University of California, Davis
Davis, California, USA
joyu@ucdavis.edu

Abstract

Conversations have an intrinsic one-to-many property, which means that multiple responses can be appropriate for the same dialog context. In task-oriented dialogs, this property leads to different valid dialog policies towards task completion. However, none of the existing task-oriented dialog generation approaches takes this property into account. We propose a Multi-Action Data Augmentation (MADA) framework to utilize the one-to-many property to generate diverse appropriate dialog responses. Specifically, we first use dialog states to summarize the dialog history, and then discover all possible mappings from every dialog state to its different valid system actions. During dialog system training, we enable the current dialog state to map to all valid system actions discovered in the previous process to create additional state-action pairs. By incorporating these additional pairs, the dialog policy learns a balanced action distribution, which further guides the dialog model to generate diverse responses. Experimental results show that the proposed framework consistently improves dialog policy diversity, and results in improved response diversity and appropriateness. Our model obtains state-of-the-art results on MultiWOZ.

Introduction

One big challenge in dialog system generation is that multiple responses can be appropriate under the same conversation context. This challenge originated from the intrinsic diversity of human conversations. Although recent progress in sequence-to-sequence (seq2seq) learning (Sutskever, Vinyals, and Le 2014) improves dialog systems performance (Serban et al. 2017; Wen et al. 2017; Lei et al. 2018). These systems still ignore this one-to-many property in conversation. Therefore, they are not able to handle diverse user behaviors in real-world settings (Li et al. 2016; Rajendran et al. 2018).

Previous studies model this one-to-many conversation property to improve *utterance-level* diversity in open-domain dialog generation (Zhao, Zhao, and Eskenazi 2017;

*This work was done during Yichi Zhang’s summer internship at University of California, Davis.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

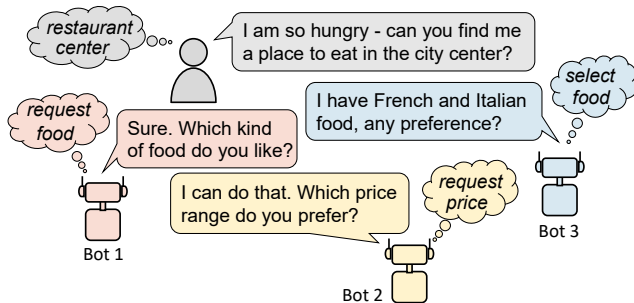


Figure 1: Multiple responses produced by different dialog policies (shown in clouds) are proper for the same context.

Zhou et al. 2017; 2018). None of previous task-oriented systems consider such one-to-many property, since they focus on task completion policies instead of language variations. However, the one-to-many phenomenon is also prevalent in task-oriented dialogs, in the form of different responding policies for the same dialog context (Fig.1). Since in collected dialog datasets each dialog context has only one reference response, the distribution of valid system actions for each dialog state rely on their occurring frequencies in the datasets which are usually highly unbalanced. Models trained on these unbalanced datasets tend to capture the most common dialog policy but ignore rarely occurred yet feasible user behaviors, which results in learning skewed and low-coverage policies.

Our goal is to address such data bias and model this one-to-many property in task-oriented dialogs to enrich dialog policy diversity, therefore building dialog systems that can generate more diverse system responses. Instead of simply learning how to map one user response to many system responses (Rajendran et al. 2018), we propose to discover the mapping from one dialog state (condensed dialog history) to multiple system actions and then generate system responses conditioned on learned actions. Since the number of unique dialog states and system actions are much smaller than the number of unique user and system responses, the mapping is more structured and easier to incorporate in learning.

Specifically, we propose a general Multi-Action Data

Augmentation (MADA) framework to achieve such mapping. We first delexicalize all utterances to reduce surface language diversity. Then we use dialog states and system actions to achieve condensed but sufficient information representation. We accumulate all mappings from dialog states to valid system actions from the entire training corpus. Finally in the dialog system training process, we force the model to not only take the ground truth system action as training sample, but also create extra training samples by including other possible system actions that are valid under that dialog state based on the state-action mapping obtained earlier. Then the learned policy is able to produce a more balanced system action distribution given a dialog context. Therefore, the dialog system can produce a set of diverse and valid system actions, which further guide the model to generate diverse and appropriate responses.

We evaluate the proposed method on MultiWOZ (Budzianowski et al. 2018), a large-scale human-human task-oriented dialog dataset covering multiple domains. We show that our data augmentation method significantly improves response generation quality on various learning models. To utilize the one-to-many mapping under the challenging multi-domain scenario, we propose Domain Aware Multi-Decoder (DAMD) network to accommodate state-action pair structure in generation. Our model obtains the new state-of-the-art results on MultiWOZ’s response generation task. Human evaluation results show that DAMD with data augmentation generates diverse and valid responses.

Related Work

The trend of building task-oriented systems is changing from training separate dialog models independently (Young et al. 2013; Wen et al. 2015; Liu and Lane 2016; Mrkšić et al. 2017) to end-to-end trainable dialog models (Zhao et al. 2017; Wen et al. 2017; Eric and Manning 2017). Specifically, Lei et al. (2018) propose a two stage seq2seq model (Sequicity) with copy mechanism (Gu et al. 2016) that completes the dialog state tracking and response generation jointly via a single seq2seq architecture. These systems achieve promising results, however, all of these models are designed for a specific domain which lacks the generalization ability to multi-domains, e.g. the recently proposed multi-domain dataset MultiWOZ (Budzianowski et al. 2018). Although several models are proposed to handle the multi-domain response generation task (Zhao, Xie, and Eskenazi 2019; Mehri, Srinivasan, and Eskenazi 2019; Chen et al. 2019), the generation quality is far from perfect, presumably due to the complex task definitions, large policy coverage and flexible language styles. We believe by modeling the one-to-many dialog property, we can improve multi-domain dialog system generation.

The one-to-many problem is more noticeable in open-domain social dialog systems since “I don’t know” can be valid response to all questions, but such response is not very useful or engaging. Therefore, previous social response generation methods attempt to model the one-to-many property by modeling responses with other meta information, such as response specificity (Zhou et al. 2017; 2018). By considering these meta information, the model

can generate social dialog response with larger diversity. However, for task-oriented dialog systems, the only work that models this one-to-many property utilizes this property to retrieve dialog system responses instead of generating response (Rajendran et al. 2018). We propose to take advantage of this one-to-many mapping property to generate more diverse dialog responses in task-oriented dialog systems. Moreover, one key advantage of our proposed framework is that the multiple actions decoded by the dialog model are interpretable and controllable. We leverage different diverse decoding methods (Li, Monroe, and Jurafsky 2016; Fan, Lewis, and Dauphin 2018; Holtzman et al. 2019) to improve the diversity of generated system actions.

Multi-Action Data Augmentation Framework

We introduce the Multi-Action Data Augmentation (MADA) framework that is generalizable to all task-oriented dialog scenarios. MADA is suitable to all dialog models that take system action supervision. It aims to increase dialog response generation diversity through learning a dialog policy that decodes a diverse set of valid system actions when given a dialog context. In MADA, we first discover the one-to-many mapping of a summarized dialog context (i.e. dialog state) to a set of system actions that are appropriate under that context. We then make the dialog model to include all the additional actions that are valid according to the one state to many system action mapping during training. In this way, the dialog policy is trained by a balanced mapping between dialog state and different system actions. Therefore, in the end the policy can generate a diverse set of system actions that are all appropriate under a give context. Such a diverse set of system actions will naturally lead to diverse system responses. Fig.2 shows an example dialog state to multiple actions mapping.

To learn this one-to-many mapping, we first need to design suitable dialog state and system action that are sufficient to represent dialog policy learning. Dialog state needs to summarize the dialog history that contains sufficient information for a dialog system to decide what actions to take next. So we define dialog state S_t at turn t to have four types of information: 1) current dialog domain, 2) belief state, 3) database search results and 4) current user action. Current dialog domain D_t is essential, because one single task can have multiple dialog domains, so the active domain is necessary to include in the dialog state representation. Belief state B_t is necessary because the belief state records slots and corresponding values informed by user in each turn, e.g. “*price=cheap, location=west*”, these slots are useful in searching database to obtain task information. Database (DB) search results DB_t also influence the next system action, because based on the data search results, the system may request for an unmentioned slot to reduce the search range. Finally current user action A_t^U can also influence the system policy, because sometimes the system need to give direct feedback to the user, such as providing a phone number when it is asked.

$$S_t \triangleq \langle D_t, B_t, DB_t, A_t^U \rangle \quad (1)$$

System action is the semantic representation of the sys-

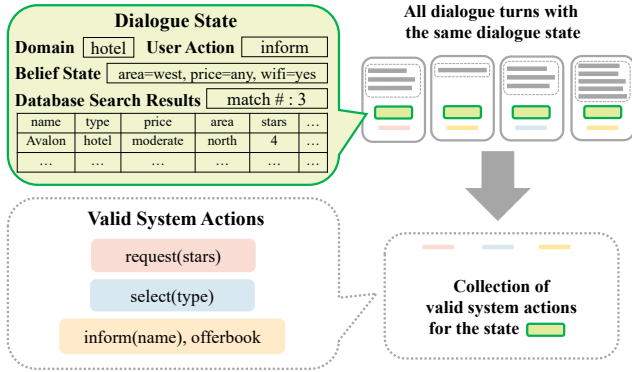


Figure 2: The overview of our Multi-Action Data Augmentation (MADA) framework. The green blocks denotes the same dialog state, and bars with different colors are different valid system actions corresponding to this state. Other valid state-action pairs are additional training data to learn the state-to-action mapping.

tem utterance. We define system action consists of dialog domains, dialog acts, and slots. One example system action is “*hotel-request(price, area)*”.

We then go through the entire training data to find system response that share the same dialog state to form all the one state to many action mappings. Finally, we introduce how to train a dialog policy that produces a balanced valid action distribution under each dialog state.

Training a dialog policy is to learn the optimal mapping from dialog states to system actions towards achieving task goals efficiently. In another way, we are learning the correct dialog actions conditioned on a dialog state:

$$\mathcal{L} = \sum_{t \in \mathcal{D}} \log P(A_t | S_t) \quad (2)$$

Due to the one-to-many property, for a specific dialog state S , there exists K different system actions $A^{(1)}, \dots, A^{(K)}$ that are valid for this state, i.e. for $i = 1, \dots, K$, $\exists t \in \mathcal{D}$ s.t. $(S_t, A_t) = (S, A^{(i)})$, and we denote the valid system action set as $\mathcal{V}(S)$. If some state-action pairs $(S, A^{(j)})$ have much lower frequency than other pairs $(S, A^{(k)})$, then the model tends to only capture the majority mappings and ignores the minority ones. So the trained dialog policy lacks diversity. This problem is also known as a general drawback of the maximum likelihood estimate over unbalance dataset (Jennrich and Schluchter 1986).

We address this issue by balancing the valid action distribution in every dialog state, S_t . Specifically, for each dialog turn t with state-action pair (S_t, A_t) , we incorporate other valid system actions under the state S_t , i.e. $A_{t'}, t' \neq t$ with $S_{t'} = S_t$, as additional training data for turn t . The new objective function is:

$$\mathcal{L}_{aug} = \sum_{t \in \mathcal{D}} \sum_{A_{t'} \in \mathcal{V}^*(S_t)} \log P(A_{t'} | S_t) \quad (3)$$

where $\mathcal{V}^*(S_t) \subseteq \mathcal{V}(S_t)$ is a subset of the valid action set $\mathcal{V}(S_t)$ of dialog state S_t . If we simply choose $\mathcal{V}^*(S_t) =$

$\mathcal{V}(S_t)$, as every $P(A|S_t)$ is optimized by exactly the same number of each valid system action corresponding to state S_t , the overall conditional probability $P(A|S)$ is optimized on a balanced set of dialog actions. ???

Our data augmentation framework over-samples training data to handle the unbalanced data problem. We choose over-sampling instead of under-sampling to make sure the dialog model can learn from all available dialogs. In practice, we can choose different $\mathcal{V}^*(S_t)$ to achieve different level of action diversity. For example, we find that for some system actions such as recommending a hotel name, a combination of other slots such as “*price*”, “*stars*”, “*parking*”, “*wifi*” etc are often informed together as additional information, which makes the number of “recommend” actions exponentially larger. However, they are semantically similar to each other. To avoid over-sampling of these actions, we group valid system actions with the same dialog act type together and uniformly sample from each group to form $\mathcal{V}^*(S_t)$. This trick improve the learning efficiency and achieves a higher action diversity over different action types. In our experiments, we find some system actions are labeled incorrectly in MultiWOZ, which makes many dialog states only have one corresponding valid system action. To address this problem, we sample $\min(K, |G|)$ actions in each group, where $K > 1$ is the predefined sample number and $|G|$ is the group size. This setting mitigates the influence of those unexpected single action groups but maintains the ability to learn from real single group action groups, e.g. rare cases in the dataset. Because large K has a negative influence on the act-level balance over small groups. We empirically set $K = 3$, as it yields the best experimental results.

As a general learning framework, MADA is applicable to any task-oriented dialog system model that takes system actions as supervision, because without the system action annotation, we would not be able to obtain state-action mappings. Our framework is suitable for all types of tasks as well. We choose the most challenging multi-domain task-oriented dialog corpus, MultiWOZ (Budzianowski et al. 2018) to validate our framework’s performance. We also designed a model, Domain Aware Multi-Decoder Network to take the full advantage of our data augmentation framework.

Domain Aware Multi-Decoder Network

We propose a Domain Aware Multi-Decoder (DAMD) network, an end-to-end model designed to handle the multi-domain response generation problem through leveraging the proposed multi-action data augmentation framework. Fig.3 shows an overview of the proposed model. There are one encoder that encodes dialog context and three decoders that decodes belief span, system action and system response respectively.

Domain-Adaptive Delexicalization

We first perform delexicalization to pre-process dialog utterances to reduce surface form language variability. Similar to Wen et al. (2017), we generate delexicalized responses with placeholders for specific slot values (see examples in Fig.3), which can be filled according to

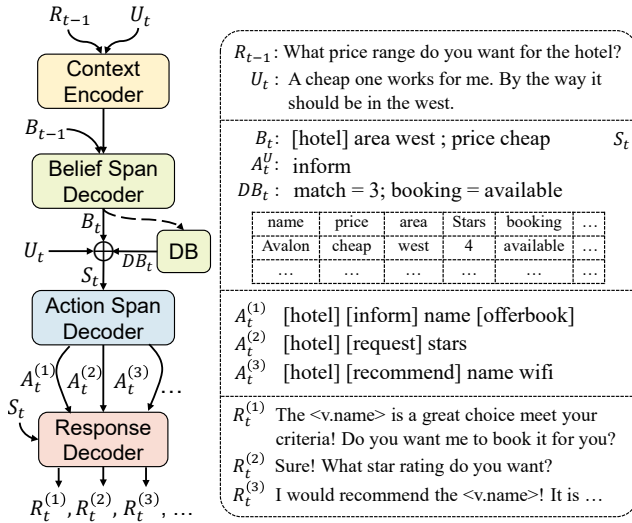


Figure 3: The overview of Domain Aware Multi-Decoder (DAMD) network. The left figure shows the information flow among all modules. The explicit inputs and outputs of each module are described on the right.

database search results afterwards. However, we find that there is a drawback in the current multi-domain delexicalization scheme (Budzianowski et al. 2018; Chen et al. 2019). Previous methods only delexicalize the same slots in different dialog domains such as *phone*, *address*, *name* etc as different tokens, e.g. $\langle restaurant.phone \rangle$ and $\langle hotel.phone \rangle$, which adds extra burdens for the system to generate these critical tokens during task completion. We propose an adaptive delexicalization scheme using one token to represent the same slot name such as $\langle v.phone \rangle$ in different dialog domains. Therefore the expressions in all relevant domains can be used to learn to generate the delexicalized value token. Since our model is domain-aware, the active domain is automatically updated based on dialog state. Therefore, there is no ambiguity in response generation process.

Belief Span Decoder

After data preprocessing, the model first learn to decode belief span. The belief span B_t of turn t is updated based on the previous belief span B_{t-1} , previous system response R_{t-1} and the current user utterance U_t through a sequence to sequence fashion:

$$B_t = \text{seq2seq}(R_{t-1}, U_t, B_{t-1}) \quad (4)$$

where the context vectors obtained by attention mechanism from the three sequences are concatenated to calculate the copy score. See Lei et al. (2018) for more details. The copy mechanism is used to copy slot names, new slot values from utterances and unchanged parts of the previous belief span. Note that the entire history utterances are not used as the context information, since all the information is already contained and summarized in belief span. But the previous response is required, since the user may have some ellipsis in

current utterance that refers to some slot values offered by system in the previous turn. The cross entropy between the generated and the ground truth belief spans are used as the loss of the belief span decoder.

In multi-domain dialog tasks, simply remembering the slot values instead of its dialog domain can lead to confusion. For example a time value can be either a reservation time in the restaurant domain or an arrival/leaving time for taxi booking. Therefore, we extend the belief span by decoding additional domain and slot tokens to address this ambiguity. An example multi-domain dialog state looks like “[*restaurant*] name *Curry Garden* time *18:00* [*taxi*] leave *20:00* destination *Kings Street*”. The active dialog domain can automatically be determined by selecting the domain that recently changed semantic slot value.

For search results DB_t , we use an one-hot vector to indicate the number of matched entities and whether the booking is available or not following Budzianowski et al. (2018).

System Action Span Decoder

The system action span decoder enables DAMD to utilize the multi-action data augmentation framework. We represent the system action as a sequence of tokens in the order of domains, acts and slots, as shown in the third text box in Fig. 3.

$$A_t = \text{seq2seq}(U_t, B_t, DB_t) \quad (5)$$

where the database search results are concatenated with the hidden state of the utterance and the belief states. We use the method described in the augmentation framework to enrich the training data. Specifically, for each system utterance in training, we find its dialog state based on annotation, which includes its dialog domain, belief state, database search result and its dialog act. Then we find this state’s appropriate system action based on the previously learned one state to many system actions obtained in our data augmentation framework. The possible state-action pairs are used to enlarge the training set.

In testing, our action decoder naturally has the ability to generate different system actions. Traditional beam search suffers from a diversity problem that the decoder tends to generate sequences with the same root ¹ (Finkel, Manning, and Ng 2006; Li, Monroe, and Jurafsky 2016), we address the issue by diversity promoting decoding techniques such as the diverse beam search (Li, Monroe, and Jurafsky 2016), top-k sampling (Fan, Lewis, and Dauphin 2018) and nucleus sampling (Holtzman et al. 2019) to further introduce dialog policies diversity.

System Response Decoder

The final step is to generate response based on the dialog state and system action, which can be formulated as:

$$R_t = \text{seq2seq}(A_t, U_t, B_t, DB_t) \quad (6)$$

where the hidden states of the belief span decoder and the action span decoder are used as B_t and A_t . Previous response

¹For example, the model is more likely to generate “[*hotel*] inform name price” together with “[*hotel*] inform name” than “[*hotel*] recommend name” by standard beam search algorithm.

decoder methods only base on system dialog act to decode sentences. Our model is trained in an end-to-end manner, where all three decoders’ loss are summed together and optimized jointly. During evaluation, different responses are generated based on different system actions.

Dataset

We evaluate our proposed framework and model on the MultiWOZ dataset (Budzianowski et al. 2018). It is a large-scale human-human task-oriented dialog dataset collected via the Wizard-of-Oz framework where one participant plays the role of the system. MultiWOZ contains conversations between a tourist and a clerk at an information center, which consists of seven domains including *hotel*, *restaurant*, *attraction*, *train*, *taxi*, *hospital* and *police*, and an additional domain *general* for some general acts such as greeting or goodbye. Each dialog in the dataset covers one to three domains, and multiple different domains might be mentioned in a single turn sometimes. Refer to Budzianowski et al. (2018) for statistics. Due to the multi-domain setting, complex ontology and flexible human expressions, developing dialog systems on MultiWOZ is extremely challenging.

Experimental Settings

Pre-processing The dataset is pre-processed through the proposed domain-adaptive delexicalization scheme as described in the previous section. The original belief state labels and system action labels are converted to the span form to train our domain-aware multi-decoder network model. The user action labels are adopted from the automatic annotations proposed by Lee et al. (2019).

Automatic Evaluation Metrics We focus on the context-to-response generation task proposed for MultiWOZ (Budzianowski et al. 2018) and follow their automatic evaluation metrics. There are four automatic metrics to evaluate the response quality - if the system provides an correct entity (**inform rate**), answers all the requested information (**success rate**), is fluent **BLEU** (Papineni et al. 2002) and a combined score **combined score** computed via $(Inform + Success) \times 0.5 + BLEU$ as an overall quality measure suggested in Mehri et al. (2019). Since our goal is to learn diversified valid actions, we introduce two additional metrics to measure the action diversity: the number of unique type of dialog acts (**act number**) and slots (**slot number**) in all generated system actions in each dialog turn. In all of our experiments we report the average score of each metric over 5 runs.

Baselines and Model Variations

We compare several model variations of our domain aware multi-decoder (DAMD) network together with other baselines on MultiWOZ.

- Seq2Seq + Attention (Budzianowski et al. 2018): a basic seq2seq model with attention (Bahdanau et al. 2014).
- Seq2Seq + Copy: a simplified version of DAMD where the belief and action span decoders are removed, which

Model & Decoding Scheme	Act #		Slot #	
	w/o	w/	w/o	w/
Single-Action Baselines				
DAMD + greedy	1.00	1.00	1.95	2.51
HDSA + fixed threshold	1.00	1.00	2.07	2.40
5-Action Generation				
DAMD + beam search	2.67	2.87	3.36	4.39
DAMD + diverse beam search	2.68	2.88	3.41	4.50
DAMD + top-k sampling	3.08	3.43	3.61	4.91
DAMD + top-p sampling	3.08	3.40	3.79	5.20
HDSA + sampled threshold	1.32	1.50	3.08	3.31
10-Action Generation				
DAMD + beam search	3.06	3.39	4.06	5.29
DAMD + diverse beam search	3.05	3.39	4.05	5.31
DAMD + top-k sampling	3.59	4.12	4.21	5.77
DAMD + top-p sampling	3.53	4.02	4.41	6.17
HDSA + sampled threshold	1.54	1.83	3.42	3.92

Table 1: Multi-action evaluation results. The “w” and “w/o” column denote with and without data augmentation respectively, and the better score between them is in bold. We report the average performance over 5 runs.

is equivalent to the copy-based seq2seq model (Gu et al. 2016).

- MD-Sequicity: a simplified version of DAMD with the action span decoder removed. We call it MD-Sequicity since it only extends the belief span to support multi-domain belief tracking comparing to the original Sequicity model (Lei et al. 2018).
- SFN + RL (Mehri et al. 2019): a seq2seq network comprised of several pre-trained dialog modules which are connected through hidden states. Reinforcement fine tuning is used additionally to train the model. SFN is similar to our model in the spirit of modeling belief state and system action jointly in an end-to-end manner, but they use binary vectors for state and action modeling and do not take advantage of copying mechanism.
- HDSA: a hierarchical disentangled self-attention network (Chen et al. 2019). A BERT-based (Devlin et al. 2018) action predictor is used to predict system actions in HDSA. Since the original multi-label classification with a fixed active threshold is not able to generate multiple actions, we alternatively samples a threshold for each dimension of the action vector independently. The actions are used to control the structure of a self-attention network afterwards for response generation, which is trained separately with the action predictor.
- DAMD: our proposed domain aware multi-decoder network. The belief state, system action and response are generated in a seq2seq manner in DAMD. We use greedy decoding for all single-sequence decoding process. When decoding multiple actions, we leverage the standard beam search algorithm and several diversity-promoted decoding schemes :
 - (1) the diverse beam search (Li et al. 2016) which adds a penalty term to intra-sibling sequences thus favors choosing hypotheses from diverse parents.

Model	Belief State Type	System Action Type Form		Inform (%)	Success (%)	BLEU	Combined Score
1. Seq2Seq + Attention (Budzianowski et al. 2018)	oracle	-	-	71.3	61.0	18.9	85.1
2. Seq2Seq + Copy	oracle	-	-	86.2	72.0	15.7	94.8
3. MD-Sequicity	oracle	-	-	86.6	71.6	16.8	95.9
4. SFN + RL (Mehri et al. 2019)	oracle	generated	one-hot	82.7	72.1	16.3	93.7
5. HDSA (Chen et al. 2019)	oracle	generated	graph	82.9	68.9	23.6	99.5
6. DAMD	oracle	generated	span	89.5	75.8	18.3	100.9
7. DAMD + multi-action data augmentation	oracle	generated	span	89.2	77.9	18.6	102.2
8. SFN + RL (Mehri et al. 2019)	oracle	oracle	one-hot	-	-	29.0	106.0
9. HDSA (Chen et al. 2019)	oracle	oracle	graph	87.9	78.0	30.4	113.4
10. DAMD + multi-action data augmentation	oracle	oracle	span	95.4	87.2	27.3	118.5
11. SFN + RL (Mehri et al. 2019)	generated	generated	one-hot	73.8	58.6	16.9	83.0
12. DAMD + multi-action data augmentation	generated	generated	span	76.3	60.4	16.6	85.0

Table 2: Comparison of response generation results on MultiWOZ. The oracle/generated denotes either using ground truth or generated results. The results are grouped according to whether and how system action is modeled.

(2) the top- k sampling algorithm (Fan et al. 2018) which samples the next word from the k most probable choices according to vocabulary distribution.

(3) the top- p sampling algorithm (Holtzman et al. 2019) which samples from the set of top possible words where their summed probability reaches a fixed value p .

Parameter Setting In our implementation of DAMD, we use a one-layer bi-directional GRU with hidden size of 100 as encoder and three standard GRUs with the same hidden size as decoders. The embedding size, vocabulary size and batch size are 50, 3,000 and 128 respectively. The combined score on development set is used as the validation check metric. We use the Adam optimizer with a initial learning rate of 0.005. The learning rate decays by half every 3 epochs if no improvement is observed on development set. Training early stops when no improvement is observed on development set for 5 epochs. For multi-action decoding, the beam size and sampling number k are the same as action number, which is 5 or 10 in our experiments. We use 0.2 as the diverse beam search penalty and $p = 0.9$ for top- p sampling. The fixed active threshold for HDSA is 0.4, and the sampling range is [0.3, 0.5] in multi-action experiments. All of the hyperparameters are selected through grid search. The code is available here².

Results and Analysis

We first evaluate whether our data augmentation framework efficiently improves dialog policy diversity. We conduct experiments of 5-action and 10-action generation, where different model variations with and without utilizing the proposed multi-action data augmentation framework are compared. The results are shown in Table 1. After applying our data augmentation, both the action and slot diversity are improved consistently, which indicates that our data augmentation framework is applicable to different models. Top- k sampling achieves the highest act-level diversity, where there are 3.43 unique dialog acts on average in five generated actions. HDSA has the worse performance and benefits less from data augmentation comparing to our proposed

domain-aware multi-decoder network, because HDSA does not decode its dialog act but perform multi-label classification. While the appropriateness of multiple actions is hard to judge by automatic evaluation (Liu et al. 2016), we leave it for human evaluation, where we also take a further step to directly evaluate the corresponding responses.

We evaluate our domain-aware multi-decoder (DAMD) network on the context-to-response generation task based on MultiWOZ. Each model generates one response for fair comparison. Experiments with ground truth belief state feed the oracle belief state as input and database search condition. Specifically in DAMD, we feed the oracle token at each decoding step of belief span to produce the oracle hidden states as input of subsequent modules. Results are shown in Table 2. The first group shows that after applying our domain-adaptive delexicalization and domain-aware belief span modeling, the task completion ability of seq2seq models becomes better. The relative lower BLEU score is potentially due to that task-promoting structures (e.g. copy) make the model focus less on learning the language surface. Our DAMD model significantly outperforms other models with different system action forms in terms of inform and success rates, which shows the superiority of our action span. While we find applying our data augmentation achieves a limited improvement on combined score (6 vs 7), which suggests learning from a balanced state-action training data can improve the robustness of model but the benefit of learning diverse policy for single response generation is hard to evaluate. Moreover, if a model has access to ground truth system action, the model further improves its task performance. Finally, we find conditioned on generated belief state greatly harm the response quality, due to the error propagation from previous decoders to the final response decoder. Note that HDSA cannot track belief state thus has no results here.

Case Study and Error Analysis

Table 3 shows an example where learning policy diversity is beneficial for task completion. Since there are still nine hotels which fit the user’s requirement, a common policy should be requesting a slot (e.g. area located) to further reduce database search range. However, the dialogs are carried

²<https://gitlab.com/ucdavisnlp/damd-multiwoz>

SYS: I also have 3 pricing options and amenity options. Could you give me some direction?	
USER: Sure. 4 star, nothing but the best, free wifi moderately priced and free parking too.	
STATE: [hotel] parking yes; pricerange moderate; stars 4; internet yes; DB-match:9	
Generated Actions w/o MADA	Generated Actions w/ MADA
[inform] area choice price	[inform] choice [request] area
[inform] area price choice type	[inform] name internet parking area [offerbook]
[inform] choice [request] area	[recommend] name [inform] choice
[inform] choice [request] area price	[recommend] name [offerbook]
[inform] choice type [request] area	[recommend] name [inform] choice [offerbook]
SYS: There are 9 moderate places in the north.	SYS: I would suggest acorn guest house. Would you like me to book you a room?
GT Action: [recommend] price name [offerbook]	
GT SYS: May I recommend acorn guest house? It is moderate and fits all your criteria. Would you like me to reserve you any rooms?	

Table 3: Our model’s example generation responses with and without data augmentation. GT denotes the ground truth. The generated action candidate closest to the ground truth action is marked in bold.

Policy Error	USER: I will be travelling from Cambridge actually and going to London Kings Cross.	USER: Yes, i would like a reservation.
	GT Action: [request] day	GT Action: [request] people
	Generated Actions: [request] leave [request] leave arrive [inform] choice [request] leave [inform] choice [request] leave arrive [inform] id arrive destination [offerbook]	Generated Actions: [offerbook] [offerbook] [general] [reqmore] [inform] food area name [offerbook] [inform] address name [offerbook] [inform] food area name [reqmore]
NLG Error	Action: [recommend] name [inform] postcode address [reqmore] SYS: I would suggest the <v.name>. The postcode is <v.postcode>. Is there anything else I can help you with today?	

Table 4: Examples of errors made by DAMD. GT denotes ground truth.

out by a large number of different crowd workers. Some workers may choose to make a direct recommendation instead. This less frequency seen policy is difficult for the system to capture without a balanced data set, as the model tend to generate only the majority *request* actions. After applying data augmentation, the *recommend* actions are also captured as a valid action.

Although better than models trained on unbalanced state-action dataset, our model still makes several types of errors shown in Table 4. The example on the left shows the model makes an error on the slot type. This is because our data augmentation method mainly focuses on improving the act-level policy diversity and the slot-level diversity is ignored. The example on the right shows an error where our system failed to collect enough information, such as number of people, before offering to make a restaurant reservation. This suggests that more prior task knowledge should be injected in the dialog model to address such issue. Moreover, besides policy level errors, there is also errors in the response generation process. In the bottom example shown in Table 5, the address information of the restaurant is missing. Such error might be caused by the generation model forgetting the distant information when the conditioned action span is too long.

Human Evaluation

Automatic metrics only validate system’s performance on one single dimension at a time. While human can provide an ultimate holistic evaluation. We conduct human evaluation to show that learning a balanced dialog policy can eventually improve the dialog system responding quality, in terms of higher appropriateness of individual responses and higher diversity among multiple responses.

In our experiments, **appropriateness** is scored on a Likert scale of 1-3 which denotes *invalid*, *ok* and *good* respectively, for each generated response. **Diversity** is scored on a Likert scale of 1-5 for all of the responses (we generate 5 responses for each model in our experiments). We suggest the judges to score according to the number of different policies in responses. We evaluate three models: DAMD without data augmentation, DAMD with data augmentation and HDSA with data augmentation. The top-*k* sampling is selected as our decoding methods since it achieves highest action diversity as shown in Table 1. We sample one hundred dialog turns and the 15 responses (five responses for each model) of each turn are scored by three judges given the dialog history.

The results are shown in Table 5. We report the average value of diversity and appropriateness, and the percentage of responses scored for each appropriateness level. With data augmentation, our model obtains a significant improvement in diversity score and achieves the best average appropriateness score as well. Due to the larger diversity, DAMD with augmentation is more likely to generate responses with better quality. However, the slightly increased invalid response percentage indicates that some invalid actions are also captured, which may due to that noisy state and action labels lead to wrong valid state-action set. We also observe our DAMD model outperforms HDSA in both diversity and appropriateness scores. This is mainly because our model considers the dialog domain information in a more effective manner and our model is able to leverage the state-action augmentation better by decoding system actions instead of performing classification. In summary, the overall results suggest that our framework can effectively improve the ability of dialog systems to generate appropriate responses with different dialog policies.

Model	Diversity	App	Good%	OK%	Invalid%
DAMD	3.12	2.50	56.5%	37.4%	6.1%
DAMD (+)	3.65	2.53	63.0%	27.1%	9.9%
HDSA (+)	2.14	2.47	57.5%	32.5%	10.0%

Table 5: Human evaluation results. Models with data augmentation are noted as (+). App denotes the average appropriateness score.

Conclusion

We focus on generating appropriate responses with higher diversity in task-oriented dialog systems, by learning a diversified dialog policy through considering the one-to-many dialog property. Specifically, we propose the Multi-Action Data Augmentation (MADA) framework to enable dialog

models to learn a more balanced state-to-action mapping. Our framework generalizes to all dialog tasks with belief state and system action annotated. We also propose a new domain aware multi-decoder (DAMD) model to leverage the proposed data augmentation framework. DAMD learns a more diverse state-to-action policy which not only achieves the state-of-the-art task success rate on the challenging MultiWOZ dataset, but also generates a set of responses that are both appropriate and diverse. In the future we plan to apply our method to help the modeling of diverse user behaviors.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Budzianowski, P.; Wen, T.-H.; Tseng, B.-H.; Casanueva, I.; Ultes, S.; Ramadan, O.; and Gašić, M. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. *arXiv preprint arXiv:1810.00278*.
- Chen, W.; Chen, J.; Qin, P.; Yan, X.; and Wang, W. Y. 2019. Semantically conditioned dialog response generation via hierarchical disentangled self-attention. *arXiv preprint arXiv:1905.12866*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Eric, M., and Manning, C. D. 2017. Key-value retrieval networks for task-oriented dialogue. *arXiv preprint arXiv:1705.05414*.
- Fan, A.; Lewis, M.; and Dauphin, Y. 2018. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.
- Finkel, J. R.; Manning, C. D.; and Ng, A. Y. 2006. Solving the problem of cascading errors: Approximate bayesian inference for linguistic annotation pipelines. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, 618–626. Association for Computational Linguistics.
- Gu, J.; Lu, Z.; Li, H.; and Li, V. O. 2016. Incorporating copying mechanism in sequence-to-sequence learning. *arXiv preprint arXiv:1603.06393*.
- Holtzman, A.; Buys, J.; Forbes, M.; and Choi, Y. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Jennrich, R. I., and Schluchter, M. D. 1986. Unbalanced repeated-measures models with structured covariance matrices. *Biometrics* 42(4):805–820.
- Lee, S.; Zhu, Q.; Takanobu, R.; Li, X.; Zhang, Y.; Zhang, Z.; Li, J.; Peng, B.; Li, X.; Huang, M.; et al. 2019. Convlab: Multi-domain end-to-end dialog system platform. *arXiv preprint arXiv:1904.08637*.
- Lei, W.; Jin, X.; Kan, M.-Y.; Ren, Z.; He, X.; and Yin, D. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1437–1447.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 110–119.
- Li, J.; Monroe, W.; and Jurafsky, D. 2016. A simple, fast diverse decoding algorithm for neural generation. *arXiv preprint arXiv:1611.08562*.
- Liu, B., and Lane, I. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. *arXiv preprint arXiv:1609.01454*.
- Liu, C.-W.; Lowe, R.; Serban, I.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2122–2132.
- Mehri, S.; Srinivasan, T.; and Eskenazi, M. 2019. Structured fusion networks for dialog. *arXiv preprint arXiv:1907.10016*.
- Mrkšić, N.; Séaghdha, D. Ó.; Wen, T.-H.; Thomson, B.; and Young, S. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1777–1788.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, 311–318. Association for Computational Linguistics.
- Rajendran, J.; Ganhotra, J.; Singh, S.; and Polymenakos, L. 2018. Learning end-to-end goal-oriented dialog with multiple answers. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3834–3843.
- Serban, I. V.; Sordoni, A.; Lowe, R.; Charlin, L.; Pineau, J.; Courville, A.; and Bengio, Y. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 3104–3112.
- Wen, T.-H.; Gasic, M.; Mrkšić, N.; Su, P.-H.; Vandyke, D.; and Young, S. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745*.
- Wen, T.-H.; Vandyke, D.; Mrkšić, N.; Gasic, M.; Barahona, L. M. R.; Su, P.-H.; Ultes, S.; and Young, S. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, 438–449.
- Young, S.; Gašić, M.; Thomson, B.; and Williams, J. D. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE* 101(5):1160–1179.
- Zhao, T.; Lu, A.; Lee, K.; and Eskenazi, M. 2017. Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability. *arXiv preprint arXiv:1706.08476*.
- Zhao, T.; Xie, K.; and Eskenazi, M. 2019. Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models. *arXiv preprint arXiv:1902.08858*.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. *arXiv preprint arXiv:1703.10960*.
- Zhou, G.; Luo, P.; Cao, R.; Lin, F.; Chen, B.; and He, Q. 2017. Mechanism-aware neural machine for dialogue response generation. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Zhou, G.; Luo, P.; Xiao, Y.; Lin, F.; Chen, B.; and He, Q. 2018. Elastic responding machine for dialog generation with dynamically mechanism selecting. In *Thirty-Second AAAI Conference on Artificial Intelligence*.