

Rethinking Action Spaces for Reinforcement Learning in End-to-end Dialog Agents with Latent Variable Models

Tiancheng Zhao¹, Kaige Xie² and Maxine Eskenazi¹

¹Language Technologies Institute, Carnegie Mellon University

²Shanghai Jiao Tong University

{tianchez, max+}@cs.cmu.edu, lightyear0117@sjtu.edu.cn

Abstract

Defining action spaces for conversational agents and optimizing their decision-making process with reinforcement learning is an enduring challenge. Common practice has been to use handcrafted dialog acts, or the output vocabulary, e.g. in neural encoder decoders, as the action spaces. Both have their own limitations. This paper proposes a novel latent action framework that treats the action spaces of an end-to-end dialog agent as latent variables and develops unsupervised methods in order to induce its own action space from the data. Comprehensive experiments are conducted examining both continuous and discrete action types and two different optimization methods based on stochastic variational inference. Results show that the proposed latent actions achieve superior empirical performance improvement over previous word-level policy gradient methods on both DealOrNoDeal and MultiWoz dialogs. Our detailed analysis also provides insights about various latent variable approaches for policy learning and can serve as a foundation for developing better latent actions in future research.¹

1 Introduction

Optimizing dialog strategies in multi-turn dialog models is the cornerstone of building dialog systems that more efficiently solve real-world challenges, e.g. providing information (Young, 2006), winning negotiations (Lewis et al., 2017), improving engagement (Li et al., 2016) etc. A classic solution employs reinforcement learning (RL) to learn a dialog policy that models the optimal action distribution conditioned on the dialog state (Williams and Young, 2007). However, since there are infinite human language possibilities, an enduring challenge has been to define what the

action space is. For traditional modular systems, the action space is defined by hand-crafted semantic representations such as dialog acts and slot-values (Raux et al., 2005; Chen et al., 2013) and the goal is to obtain a dialog policy that chooses the best hand-crafted action at each dialog turn. But it is limited because it can only handle simple domains whose entire action space can be captured by hand-crafted representations (Walker, 2000; Su et al., 2017). This cripples a system’s ability to handle conversations in complex domains.

Conversely, end-to-end (E2E) dialog systems have removed this limit by directly learning a response generation model conditioned on the dialog context using neural networks (Vinyals and Le, 2015; Sordoni et al., 2015). To apply RL to E2E systems, the action space is typically defined as the entire vocabulary; every response output word is considered to be an action selection step (Li et al., 2016), which we denote as the word-level RL. Word-level RL, however, has been shown to have several major limitations in learning dialog strategies. The foremost one is that direct application of word-level RL leads to degenerate behavior: the response decoder deviates from human language and generates utterances that are incomprehensible (Lewis et al., 2017; Das et al., 2017; Kottur et al., 2017). A second issue is that since a multi-turn dialog can easily span hundreds of words, word-level RL suffers from credit assignment over a long horizon, leading to slow and sub-optimal convergence (Kaelbling et al., 1996; He et al., 2018).

This paper proposes Latent Action Reinforcement Learning (LaRL), a novel framework that overcomes the limitations of word-level RL for E2E dialog models, marrying the benefits of a traditional modular approach in an unsupervised manner. The key idea is to develop E2E models that can invent their own discourse-level ac-

¹Data and code are available at <https://github.com/snakeztc/NeuralDialog-LaRL>

tions. These actions must be expressive enough to capture response semantics in complex domains (i.e. have the capacity to represent a large number of actions), thus decoupling the discourse-level decision-making process from natural language generation. Then any RL technique can be applied to this induced action space in the place of word-level output. We propose a flexible latent variable dialog framework and investigate several approaches to inducing latent action space from natural conversational data. We further propose (1) a novel training objective that outperforms the typical evidence lower bound used in dialog generation and (2) an attention mechanism for integrating discrete latent variables in the decoder to better model long responses.

We test this on two datasets, DealOrNoDeal (Lewis et al., 2017) and MultiWoz (Budzianowski et al., 2018), to answer two key questions: (1) what are the advantages of LaRL over Word-level RL and (2) what effective methods can induce this latent action space. Results show that LaRL is significantly more effective than word-level RL for learning dialog policies and it does not lead to incomprehensible language generation. Our models achieve 18.2% absolute improvement over the previous state-of-the-art on MultiWoz and discover novel and diverse negotiation strategies on DealOrNoDeal. Besides strong empirical improvement, our model analysis reveals novel insights, e.g. it is crucial to reduce the exposure bias in the latent action space and discrete latent actions are more suitable than continuous ones to serve as action spaces for RL dialog agents.

2 Related Work

Prior RL research in modular dialog management has focused on policy optimization over hand-crafted action spaces in task-oriented domains (Walker, 2000; Young et al., 2007). A dialog manager is formulated as a Partially Observable Markov Decision Process (POMDP) (Young et al., 2013), where the dialog state is estimated via dialog state tracking models from the raw dialog context (Lee, 2013; Henderson et al., 2014; Ren et al., 2018). RL techniques are then used to find the optimal dialog policy (Gasic and Young, 2014; Su et al., 2017; Williams et al., 2017). Recent deep-learning modular dialog models have also explored joint optimization over dialog pol-

icy and state tracking to achieve stronger performance (Wen et al., 2016; Zhao and Eskenazi, 2016; Liu and Lane, 2017).

A related line of work is reinforcement learning for E2E dialog systems. Due to the flexibility of encoder-decoder dialog models, prior work has applied reinforcement learning to more complex domains and achieved higher dialog-level rewards, such as open-domain chatting (Li et al., 2016; Serban et al., 2017a), negotiation (Lewis et al., 2017), visual dialogs (Das et al., 2017), grounded dialog (Mordatch and Abbeel, 2017) etc. As discussed in Section 1, these methods consider the output vocabulary at every decoding step to be the action space; they suffer from limitations such as deviation from natural language and sub-optimal convergence.

Finally, research in latent variable dialog models is closely related to our work, which strives to learn meaningful latent variables for E2E dialog systems. Prior work has shown that learning with latent variables leads to benefits like diverse response decoding (Serban et al., 2017b; Zhao et al., 2017; Cao and Clark, 2017), interpretable decision-making (Wen et al., 2017; Zhao et al., 2018) and zero-shot domain transfer (Zhao and Eskenazi, 2018). Also, driven by similar motivations of this work, prior studies have explored to utilize a coarse discrete node, either handcrafted or learned, to decouple the word generation process from dialog policy in E2E systems for better dialog policy (He et al., 2018; Yarats and Lewis, 2017). Our work differs from prior work for two reasons: (1) latent action in previous work is only auxiliary, small-scale and mostly learned in a supervised or semi-supervised setting. This paper focuses on unsupervised learning of latent variables and learns variables that are expressive enough to capture the entire action space by itself. (2) to our best knowledge, our work is the first comprehensive study of the use of latent variables for RL policy optimization in dialog systems.

3 Baseline Approach

E2E response generation can be treated as a conditional language generation task, which uses neural encoder-decoders (Cho et al., 2014) to model the conditional distribution $p(\mathbf{x}|\mathbf{c})$ where \mathbf{c} is the observed dialog context and \mathbf{x} is the system's response to the context. The format of the dialog context is domain dependent. It can vary from tex-




Figure 1: High-level comparison between word-level and latent-action reinforcement learning in a sample multi-turn dialog. The decoder network generates the response given the latent code \mathbf{z} . Dashed line denotes places where policy gradients from task rewards are applied to the model.

tual raw dialog history (Vinyals and Le, 2015) to visual and textual context (Das et al., 2017). Training with RL usually has 2 steps: supervised pre-training and policy gradient reinforcement learning (Williams and Zweig, 2016; Dhingra et al., 2017; Li et al., 2016). Specifically, the supervised learning step maximizes the log likelihood on the training dialogs, where θ is the model parameter:

$$\mathcal{L}_{SL}(\theta) = \mathbb{E}_{\mathbf{x}, \mathbf{c}}[\log p_{\theta}(\mathbf{x}|\mathbf{c})] \quad (1)$$

Then the following RL step uses policy gradients, e.g. the REINFORCE algorithm (Williams, 1992) to update the model parameters with respect to task-dependent goals. We assume that we have an environment that the dialog agent can interact with and that there is a turn-level reward r_t at every turn t of the dialog. We can then write the expected discounted return under a dialog model θ as $J(\theta) = \mathbb{E}[\sum_0^T \gamma^t r_t]$, where $\gamma \in [0, 1]$ is the discounting factor and T is the length of the dialog. Often a baseline function b is used to reduce the variance of the policy gradient (Greensmith et al., 2004), leading to $R_t = \sum_{k=0}^{T-t} \gamma^k (r_{t+k} - b)$.

Word-level Reinforcement Learning: as shown in Figure 1, the baseline approach treats every output word as an action step and its policy gradient is:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\theta} \left[\sum_{t=0}^T \sum_{j=0}^{U_t} R_{tj} \nabla_{\theta} \log p_{\theta}(w_{tj} | w_{<tj}, \mathbf{c}_t) \right] \quad (2)$$

where U_t is the number of tokens in the response at turn t and j is the word index in the response. It is evident that Eq 2 has a very large action space, i.e. $|V|$ and a long learning horizon, i.e. TU . Prior work has found that the direct application of Eq 2 leads to divergence of the decoder. The common solution is to alternate with supervised learning with Eq 2 at a certain ratio (Lewis et al., 2017). We denote this ratio as RL:SL=A:B, which means

for every A policy gradient updates, we run B supervised learning updates. We use RL:SL=off for the case where only policy gradients are used and no supervised learning is involved.

4 Latent Action Reinforcement Learning

We now describe the proposed LaRL framework. As shown in Figure 1, a latent variable \mathbf{z} is introduced in the response generation process. The conditional distribution is factorized into $p(\mathbf{x}|\mathbf{c}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z}|\mathbf{c})$ and the generative story is: (1) given a dialog context \mathbf{c} we first sample a latent action \mathbf{z} from $p_{\theta_e}(\mathbf{z}|\mathbf{c})$ and (2) generate the response by sampling \mathbf{x} based on \mathbf{z} via $p_{\theta_d}(\mathbf{x}|\mathbf{z})$, where p_{θ_e} is the dialog encoder network and p_{θ_d} is the response decoder network. Given the above setup, LaRL treats the latent variable \mathbf{z} as its action space instead of outputting words in response \mathbf{x} . We can now apply REINFORCE in the latent action space:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\theta} \left[\sum_{t=0}^T R_t \log p_{\theta}(\mathbf{z}|\mathbf{c}_t) \right] \quad (3)$$

Compared to Eq 2, LaRL differs by:

- Shortens the horizon from TU to T .
- Latent action space is designed to be low-dimensional, much smaller than V .
- The policy gradient only updates the encoder θ_e and the decoder θ_d stays intact.

These properties reduce the difficulties for dialog policy optimization and decouple high-level decision-making from natural language generation. The p_{θ_e} are responsible for choosing the best latent action given a context \mathbf{c} while p_{θ_d} is only responsible for transforming \mathbf{z} into the surface-form words. Our formulation also provides a flexible framework for experimenting with various types of model learning methods. In this paper, we focus on two key aspects: the type of latent variable

\mathbf{z} and optimization methods for learning \mathbf{z} in the supervised pre-training step.

4.1 Types of Latent Actions

Two types of latent variables have been used in previous research: continuous isotropic Gaussian distribution (Serban et al., 2017b) and multivariate categorical distribution (Zhao et al., 2018). These two types are both compatible with our LaRL framework and can be defined as follows:

Gaussian Latent Actions follow M dimensional multivariate Gaussian distribution with a diagonal covariance matrix, i.e. $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2 \mathbf{I})$. Let the encoder p_{θ_e} consist of two parts: a context encoder \mathcal{F} , a neural network that encodes the dialog context \mathbf{c} into a vector representation \mathbf{h} , and a feed forward network π that projects \mathbf{h} into $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$. The process is defined as follows:

$$\mathbf{h} = \mathcal{F}(\mathbf{c}) \quad (4)$$

$$\begin{bmatrix} \boldsymbol{\mu} \\ \log(\boldsymbol{\sigma}^2) \end{bmatrix} = \pi(\mathbf{h}) \quad (5)$$

$$p(\mathbf{x}|\mathbf{z}) = p_{\theta_d}(\mathbf{z}) \quad \mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2 \mathbf{I}) \quad (6)$$

where the sampled \mathbf{z} is used as the initial state of the decoder for response generation. Also we use $p_{\theta}(\mathbf{z}|\mathbf{c}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\sigma}^2 \mathbf{I})$ to compute the policy gradient update in Eq 3.

Categorical Latent Actions are M independent K -way categorical random variables. Each \mathbf{z}_m has its own token embeddings to map latent symbols into vector space $\mathbf{E}_m \in \mathbb{R}^{K \times D}$ where $m \in [1, M]$ and D is the embedding size. Thus M latent actions can represent exponentially, K^M , unique combinations, making it expressive enough to model dialog acts in complex domains. Similar to Gaussian Latent Actions, we have

$$\mathbf{h} = \mathcal{F}(\mathbf{c}) \quad (7)$$

$$p(Z_m|\mathbf{c}) = \text{softmax}(\pi_m(\mathbf{h})) \quad (8)$$

$$p(\mathbf{x}|\mathbf{z}) = p_{\theta_d}(\mathbf{E}_{1:M}(\mathbf{z}_{1:M})) \quad \mathbf{z}_m \sim p(Z_m|\mathbf{c}) \quad (9)$$

For the computing policy gradient in Eq 3, we have $p_{\theta}(\mathbf{z}|\mathbf{c}) = \prod_{m=1}^M p(Z_m = \mathbf{z}_m|\mathbf{c})$

Unlike Gaussian latent actions, a matrix $\mathbb{R}^{M \times D}$ comes after the embedding layers $\mathbf{E}_{1:M}(\mathbf{z}_{1:M})$, whereas the decoder’s initial state is a vector of size \mathbb{R}^D . Previous work integrated this matrix with the decoder by summing over the latent embeddings, i.e. $\mathbf{x} = p_{\theta_d}(\sum_1^M \mathbf{E}_m(\mathbf{z}_m))$, denoted as **Summation Fusion** for later discussion (Zhao

et al., 2018). A limitation of this method is that it could lose fine-grained order information in each latent dimension and have issues with long responses that involve multiple dialog acts. Therefore, we propose a novel method, **Attention Fusion**, to combine categorical latent actions with the decoder. We apply the attention mechanism (Luong et al., 2015) over latent actions as the following. Let i be the step index during decoding. Then we have:

$$\alpha_{mi} = \text{softmax}(\mathbf{h}_i^T \mathbf{W}_a \mathbf{E}_m(\mathbf{z}_m)) \quad (10)$$

$$\mathbf{c}_i = \sum_{m=1}^M \alpha_{mi} \mathbf{E}_m(\mathbf{z}_m) \quad (11)$$

$$\tilde{\mathbf{h}}_i = \tanh(\mathbf{W}_s \begin{bmatrix} \mathbf{h}_i \\ \mathbf{c}_i \end{bmatrix}) \quad (12)$$

$$p(w_i|\mathbf{h}_i, \mathbf{c}_i) = \text{softmax}(\mathbf{W}_o \tilde{\mathbf{h}}_i) \quad (13)$$

The decoder’s next state is updated by $\mathbf{h}_{i+1} = \text{RNN}(\mathbf{h}_i, w_{i+1}, \tilde{\mathbf{h}}_i)$ and \mathbf{h}_0 is computed via summation-fusion. Thus attention fusion lets the decoder focus on different latent dimensions at each generation step.

4.2 Optimization Approaches

Full ELBO: Now given a training dataset $\{\mathbf{x}, \mathbf{c}\}$, our base optimization method is via stochastic variational inference by maximizing the evidence lowerbound (ELBO), a lowerbound on the data log likelihood:

$$\mathcal{L}_{full}(\theta) = p_{q(\mathbf{z}|\mathbf{x}, \mathbf{c})}(\mathbf{x}|\mathbf{z}) - \text{D}_{\text{KL}}[q(\mathbf{z}|\mathbf{x}, \mathbf{c})||p(\mathbf{z}|\mathbf{c})] \quad (14)$$

where $q_{\gamma}(\mathbf{z}|\mathbf{x}, \mathbf{c})$ is a neural network that is trained to approximate the posterior distribution $q(\mathbf{z}|\mathbf{x}, \mathbf{c})$ and $p(\mathbf{z}|\mathbf{c})$ and $p(\mathbf{x}|\mathbf{z})$ are achieved by \mathcal{F} , π and p_{θ_d} . For Gaussian latent actions, we use the reparametrization trick (Kingma and Welling, 2013) to backpropagate through Gaussian latent actions and the Gumbel-Softmax (Jang et al., 2016) to backpropagate through categorical latent actions.

Lite ELBO: a major limitation is that Full ELBO can suffer from exposure bias at latent space, i.e. the decoder only sees \mathbf{z} sampled from $q(\mathbf{z}|\mathbf{x}, \mathbf{c})$ and never experiences \mathbf{z} sampled from $p_{\theta}(\mathbf{z}|\mathbf{c})$, which is always used at testing time. Therefore, in this paper, we propose a simplified ELBO for encoder-decoder models with stochastic latent variables:

$$\mathcal{L}_{lite}(\theta) = p_{p(\mathbf{z}|\mathbf{c})}(\mathbf{x}|\mathbf{z}) - \beta \text{D}_{\text{KL}}[p(\mathbf{z}|\mathbf{c})||p(\mathbf{z})] \quad (15)$$

Essentially this simplified objective sets the posterior network the same as our encoder, i.e. $q_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{c}) = p_{\theta_e}(\mathbf{z}|\mathbf{c})$, which makes the KL term in Eq 14 zero and removes the issue of exposure bias. But this leaves the latent spaces unregularized and our experiments show that if we only maximize $p_{p(\mathbf{z}|\mathbf{c})}(\mathbf{x}|\mathbf{z})$ there is overfitting. For this, we add the additional regularization term $\beta \text{D}_{\text{KL}}[p(\mathbf{z}|\mathbf{c})||p(\mathbf{z})]$ that encourages the posterior be similar to certain prior distributions and β is a hyper-parameter between 0 and 1. We set the $p(\mathbf{z})$ for categorical latent actions to be uniform, i.e. $p(\mathbf{z}) = 1/K$, and set the prior for Gaussian latent actions to be $\mathcal{N}(\mathbf{0}, \mathbf{I})$, which we will show that are effective.

5 Experiment Settings

5.1 DealOrNoDeal Corpus and RL Setup

DealOrNoDeal is a negotiation dataset that contains 5805 dialogs based on 2236 unique scenarios (Lewis et al., 2017). We hold out 252 scenarios for testing environment and randomly sample 400 scenarios from the training set for validation. The results are evaluated from 4 perspectives: Perplexity (PPL), Reward, Agree and Diversity. PPL helps us to identify which model produces the most human-like responses, while Reward and Agree evaluate the model’s negotiation strength. Diversity indicates whether the model discovers a novel discourse-level strategy or just repeats dull responses to compromise with the opponent. We closely follow the original paper and use the same reward function and baseline calculation. At last, to have a fair comparison, all the compared models shared the identical judge model and user simulator, which are a standard hierarchical encoder-decoder model trained with Maximum Likelihood Estimation (MLE).

5.2 Multi-Woz Corpus and Novel RL Setup

Multi-Woz is a slot-filling dataset that contains 10438 dialogs on 6 different domains. 8438 dialogs are for training and 1000 each are for validation and testing. Since no prior user simulator exists for this dataset, for a fair comparison with the previous state-of-the-art we focus on the Dialog-Context-to-Text Generation task proposed in (Budzianowski et al., 2018). This task assumes that the model has access to the ground-truth dialog belief state and is asked to generate the next response at every system turn in a di-

alog. The results are evaluated from 3 perspectives: BLEU, Inform Rate and Success Rate. The BLEU score checks the response-level lexical similarity, while Inform and Success Rate measure whether the model gives recommendations and provides all the requested information at dialog-level. Current state-of-the-art results struggle in this task and MLE models only achieve 60% success (Budzianowski et al., 2018). To transform this task into an RL task, we propose a novel extension to the original task as follows:

1. For each RL episode, randomly sample a dialog from the training set
2. Run the model on every system turn, and do not alter the original dialog context at every turn given the generated responses.
3. Compute Success Rate based on the generated responses in this dialog.
4. Compute policy gradient using Eq 3 and update the parameters.

This setup creates a variant RL problem that is similar to the Contextual Bandits (Langford and Zhang, 2008), where the goal is to adjust its parameters to generate responses that yield better Success Rate. Our results show that this problem is challenging and that word-level RL falls short.

5.3 Language Constrained Reward (LCR) curve for Evaluation

It is challenging to quantify the performance of RL-based neural generation systems because it is possible for a model to achieve high task reward and yet not generate human language (Das et al., 2017). Therefore, we propose a novel measure, the Language Constrained Reward (LCR) curve as an additional robust measure. The basic idea is to use an ROC-style curve to visualize the tradeoff between achieving higher reward and being faithful to human language. Specifically, at each checkpoint i over the course of RL training, we record two measures: (1) the PPL of a given model on the test data $p_i = \text{PPL}(\theta_i)$ and (2) this model’s average cumulative task reward in the test environment R_i^t . After RL training is complete, we create a 2D plot where the x-axis is the maximum PPL allowed, and the y-axis is the best achievable reward within the PPL budget in the testing environments:

$$y = \max_i R_i^t \quad \text{subject to} \quad p_i < x \quad (16)$$

As a result, a perfect model should lie in the upper left corner whereas a model that sacrifices language quality for higher reward will lie in the lower right corner. Our results will show that the LCR curve is an informative and robust measure for model comparison.

6 Results: Latent Actions or Words?

We have created 6 different variations of latent action dialog models under our LaRL framework. To demonstrate the advantages of LaRL, during

Model	Var Type	Loss	Integration
Gauss	Gaussian	\mathcal{L}_{full}	/
Cat	Categorical	\mathcal{L}_{full}	sum
AttnCat	Categorical	\mathcal{L}_{full}	attn
LiteGauss	Gaussian	\mathcal{L}_{lite}	/
LiteCat	Categorical	\mathcal{L}_{lite}	sum
LiteAttnCat	Categorical	\mathcal{L}_{lite}	attn

Table 1: All proposed variations of LaRL models.

the RL training step, we set RL:SL=off for all latent action models, while the baseline word-level RL models are free to tune RL:SL for best performance. For latent variable models, their perplexity is estimated via Monte Carlo $p(\mathbf{x}|\mathbf{c}) \approx \mathbb{E}_{p(\mathbf{z}|\mathbf{c})}[p(\mathbf{x}|\mathbf{z})p(\mathbf{z}|\mathbf{c})]$. For the sake of clarity, this section only compares the best performing latent action models to the best performing word-level models and focuses on the differences between them. A detailed comparison of the 6 latent space configurations is addressed in Section 7.

6.1 DealOrNoDeal

The baseline system is a hierarchical recurrent encoder-decoder (HRED) model (Serban et al., 2016) that is tuned to reproduce results from (Lewis et al., 2017). Word-level RL is then used to fine-tune the pre-trained model with RL:SL=4:1. On the other hand, the best performing latent action model is LiteCat. Best models are chosen based on performance on the validation environment.

The results are summarized in Table 2 and Figure 2 shows the LCR curves for the baseline with the two best models plus LiteAttnCat and baseline without RL:SL. From Table 2, it appears that the word-level RL baseline performs better than LiteCat in terms of rewards. However, Figure 2 shows that the two LaRL models achieve strong task rewards with a much smaller performance drop in language quality (PPL), whereas the word-level

	PPL	Reward	Agree%	Diversity
Baseline	5.23	3.75	59	109
LiteCat	5.35	2.65	41	58
Baseline +RL	8.23	7.61	86	5
LiteCat +RL	6.14	7.27	87	202

Table 2: Results on DealOrNoDeal. Diversity is measured by the number of unique responses the model used in all scenarios from the test data.

model can only increase its task rewards by deviating significantly from natural language.




Figure 2: LCR curves on DealOrNoDeal dataset.

Closer analysis shows the word-level baseline severely overfits to the user simulator. The caveat is that the word-level models have in fact discovered a loophole in the simulator by insisting on 'hat' and 'ball' several times and the user model eventually yields to agree to the deal. This is reflected in the diversity measure, which is the number of unique responses that a model uses in all 200 testing scenarios. As shown in Figure 3, after RL training, the diversity of the baseline model drops to only 5. It is surprising that the agent can achieve high reward with a well-trained HRED user simulator using only 5 unique utterances. On the contrary, LiteCat increases its response diversity after RL training from 58 to 202, suggesting that LiteCat discovers novel discourse-level strategies in order to win the negotiation instead of exploiting local loopholes in the same user simulator. Our qualitative analysis confirms this when we observe that our LiteCat model is able to use multiple strategies in negotiation, e.g. elicit preference question, request different offers, insist on key objects etc. See supplementary material for example conversations.




Figure 3: Response diversity and task reward learning curve over the course of RL training for both word RL:SL=4:1 (left) and LiteCat (right).

6.2 MultiWoz

For MultiWoz, we reproduce results from (Budzianowski et al., 2018) as the baseline. After RL training, the best LaRL model is LiteAttnCat and the best word-level model is word RL:SL=off. Table 3 shows that LiteAttnCat is on

	PPL	BLEU	Inform	Success
Human	/	/	90%	82.3%
Baseline	3.98	18.9	71.33%	60.96%
LiteAttnCat	4.05	19.1	67.98%	57.36%
Baseline	17.11	1.4	80.5%	79.07%
+RL				
LiteAttnCat	5.22	12.8	82.78%	79.2%
+RL				

Table 3: Main results on MultiWoz test set. RL models are chosen based on performance on the validation set.

par with the baseline in the supervised learning step, showing that multivariate categorical latent variables alone are powerful enough to match with continuous hidden representations for modeling dialog actions. For performance after RL training, LiteAttnCat achieves near-human performance in terms of success rate and inform rate, obtaining 18.24% absolute improvement over the MLE-based state-of-the-art (Budzianowski et al., 2018). More importantly, perplexity only slightly increases from 4.05 to 5.22. On the other hand, the word-level RL’s success rate also improves to 79%, but the generated responses completely deviate from natural language, increasing perplexity from 3.98 to 17.11 and dropping BLEU from 18.9

to 1.4.




Figure 4: LCR curves on the MultiWoz dataset.

Figure 4 shows the LCR curves for MultiWoz, with a trend similar to the previous section: the word-level models can only achieve task reward improvement by sacrificing their response decoder PPL. Figure 4 also shows the LCR curve for the baseline trained with RL:SL=100:1, hoping that supervised learning can force the model to conform to natural language. While PPL and BLEU are indeed improved, it also limits final reward performance. The latent-level models, on the contrary, do not suffer from this tradeoff. We also observe that LiteAttnCat consistently outperforms LiteCat on MultiWoz, confirming the effectiveness of Attention Fusion for handling long dialog responses with multiple entities and dialog acts. Lastly, Table 4 qualitatively exhibits the generation differences between the two approaches. The RL:SL=off model learns to continuously output entities to fool the evaluation script for high success rate, whereas LiteCatAttn learns to give more information while maintaining the language quality.

Context	Sys I have [value_count] trains matching your request . Is there a specific day and time you would like to travel? Usr I would like to leave on [value_day] and arrive by [value_time].
Model	Generated Response
word RL:SL=off	[train_id] is leaving [value_place] on [value_day] on [value_day] on [train_id] [train_id] [value_count] [train_id] leaving ...
word RL:SL=100	[train_id] leaves at [value_time] . would you like me to book you a ticket ?
LiteAttnCat	[train_id] leaves [value_place] at [value_time] and arrives in [value_place] at [value_time]. Would you like me to book that for you ?

Table 4: Example responses from baselines and LiteCatAttn on MultiWoz.

7 Model Analysis

We compare the 6 variants of latent action models on DealOrNoDeal and MultiWoz. Table 5

Deal	PPL	Reward	Agree%	Diversity
Baseline	3.23	3.75	59	109
Gauss	110K	2.71	43	176
LiteGauss	5.35	4.48	65	91
Cat	80.41	3.9	62	115
AttnCat	118.3	3.23	51	145
LiteCat	5.35	2.67	41	58
LiteAttnCat	5.25	3.69	52	75
MultiWoz	PPL	BLEU	Inform%	Succ%
Baseline	3.98	18.9	71.33	60.96
Gauss	712.3	7.54	60.5	23.0
LiteGauss	4.06	19.3	56.46	48.06
Cat	7.07	13.7	54.15	42.04
AttnCat	12.01	12.6	63.9	45.8
LiteCat	4.10	19.1	61.56	49.15
LiteAttnCat	4.05	19.1	67.97	57.36

Table 5: Comparison of 6 model variants with only supervised learning training.

shows performance of the models that are pre-trained only with supervised learning. Figure 5 shows LCR curves for the 3 models pre-trained with \mathcal{L}_{lite} and fine-tuned with policy gradient reinforcement learning. The following are the main

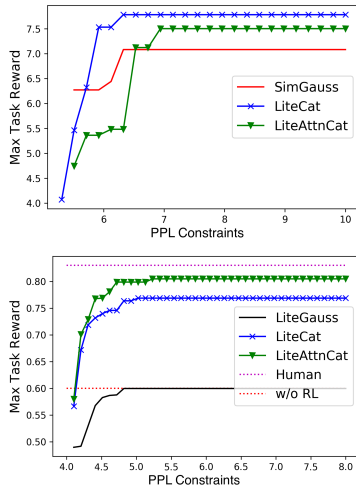


Figure 5: LCR curves on DealOrNoDeal and MultiWoz. Models with \mathcal{L}_{full} are not included because their PPLs are too poor to compare to the Lite models.

findings based on these results.

\mathcal{L}_{lite} **outperforms** \mathcal{L}_{full} **as a pre-train objective.** Table 5 shows that models with \mathcal{L}_{full} fall behind their Lite counterparts on PPL and

BLEU. We attribute this to the exposure bias in the latent space, i.e. the decoder is not trained to consider the discrepancy between the posterior network and actual dialog policy network. Meanwhile, the full models tend to enjoy higher diversity at pre-training, which agrees with the diversity-promoting effect observed in prior research (Zhao et al., 2017). However, our previous discussion on Figure 3 shows that Lite models are able to increase their response diversity in order to win more in negotiation through RL training. This is fundamentally different from diversity in pre-training, since diversity in LaRL is optimized to improve task reward, rather than to better model the original data distribution. Table 6 shows the

β	0.0	0.01	β	0.0	0.01
LiteCat	4.23	7.27	LiteGauss	4.83	6.67

Table 6: Best rewards in test environments on DealOrNoDeal with various β .

importance of latent space regularization. When β is 0, both LiteCat and LiteGauss reach suboptimal policies with final reward that are much smaller than the regularized versions ($\beta = 0.01$). The reason behind this is that the unregularized pre-trained policy has very low entropy, which prohibits sufficient exploration in the RL stage.

Categorical latent actions outperform Gaussian latent actions. Models with discrete actions consistently outperform models with Gaussian ones. This is surprising since continuously distributed representations are a key reason for the success of deep learning in natural language processing. Our finding suggests that (1) multivariate categorical distributions are powerful enough to model complex natural dialog responses semantics, and can achieve on par results with Gaussian or non-stochastic continuous representations. (2) categorical variables are a better choice to serve as action spaces for reinforcement learning. Figure 5 shows that Lite(Attn)Cat easily achieves strong rewards while LiteGauss struggles to improve its reward. Also, applying REINFORCE on Gaussian latent actions is unstable and often leads to model divergence. We suspect the reason for this is the unbounded nature of continuous latent space: RL exploration in the continuous space may lead to areas in the manifold that are not covered in supervised training, which causes undefined decoder behavior given z in these unknown areas.

8 Conclusion and Future Work

In conclusion, this paper proposes a latent variable action space for RL in E2E dialog agents. We present a general framework with a regularized ELBO objective and attention fusion for discrete variables. The methods are assessed on two dialog tasks and analyzed using the proposed LCR curve. Results show our models achieve superior performance and create a new state-of-the-art success rate on MultiWoz. Extensive analyses enable us to gain insight on how to properly train latent variables that can serve as the action spaces for dialog agents. This work is situated in the approach concerning practical latent variables in dialog agents, being able to create action abstraction in an unsupervised manner. We believe that our findings are a basic first step in this promising research direction.

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. 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*, pages 5016–5026.
- Kris Cao and Stephen Clark. 2017. Latent variable dialogue models and their diversity. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 182–187.
- Yun-Nung Chen, William Yang Wang, and Alexander I Rudnicky. 2013. Unsupervised induction and filling of semantic slots for spoken dialogue systems using frame-semantic parsing. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on*, pages 120–125. IEEE.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Abhishek Das, Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. 2017. Learning cooperative visual dialog agents with deep reinforcement learning. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2970–2979. IEEE.
- Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 484–495.
- Milica Gasic and Steve Young. 2014. Gaussian processes for pomdp-based dialogue manager optimization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(1):28–40.
- Evan Greensmith, Peter L Bartlett, and Jonathan Baxter. 2004. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(Nov):1471–1530.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling strategy and generation in negotiation dialogues. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2333–2343.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 292–299.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Satwik Kottur, José Moura, Stefan Lee, and Dhruv Batra. 2017. Natural language does not emerge naturally in multi-agent dialog. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2962–2967.
- John Langford and Tong Zhang. 2008. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in neural information processing systems*, pages 817–824.
- Sungjin Lee. 2013. Structured discriminative model for dialog state tracking. In *Proceedings of the SIGDIAL 2013 Conference*, pages 442–451.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.

- Bing Liu and Ian Lane. 2017. An end-to-end trainable neural network model with belief tracking for task-oriented dialog. *Proc. Interspeech 2017*, pages 2506–2510.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Igor Mordatch and Pieter Abbeel. 2017. Emergence of grounded compositional language in multi-agent populations. *arXiv preprint arXiv:1703.04908*.
- Antoine Raux, Brian Langner, Dan Bohus, Alan W Black, and Maxine Eskenazi. 2005. Lets go public! taking a spoken dialog system to the real world. In *Proc. of Interspeech 2005*. Citeseer.
- Liliang Ren, Kaige Xie, Lu Chen, and Kai Yu. 2018. Towards universal dialogue state tracking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2780–2786.
- Iulian V Serban, Chinnadhurai Sankar, Mathieu Germain, Saizheng Zhang, Zhouhan Lin, Sandeep Subramanian, Taesup Kim, Michael Pieper, Sarath Chandar, Nan Rosemary Ke, et al. 2017a. A deep reinforcement learning chatbot. *arXiv preprint arXiv:1709.02349*.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017b. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205.
- Pei-Hao Su, Paweł Budzianowski, Stefan Ultes, Milica Gasic, and Steve Young. 2017. Sample-efficient actor-critic reinforcement learning with supervised data for dialogue management. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 147–157.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Marilyn A. Walker. 2000. An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research*, pages 387–416.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve Young. 2017. Latent intention dialogue models. In *International Conference on Machine Learning*, pages 3732–3741.
- Jason D Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. 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)*, volume 1, pages 665–677.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- Jason D Williams and Geoffrey Zweig. 2016. End-to-end lstm-based dialog control optimized with supervised and reinforcement learning. *arXiv preprint arXiv:1606.01269*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489.
- Denis Yarats and Mike Lewis. 2017. Hierarchical text generation and planning for strategic dialogue. *arXiv preprint arXiv:1712.05846*.
- Stephanie Young, Jost Schatzmann, Karl Weilhammer, and Hui Ye. 2007. The hidden information state approach to dialog management. In *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, volume 4, pages IV–149. IEEE.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Steve J Young. 2006. Using pomdps for dialog management. In *SLT*, pages 8–13.

Tiancheng Zhao and Maxine Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 1.

Tiancheng Zhao and Maxine Eskenazi. 2018. Zero-shot dialog generation with cross-domain latent actions. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 1–10.

Tiancheng Zhao, Kyusong Lee, and Maxine Eskenazi. 2018. Unsupervised discrete sentence representation learning for interpretable neural dialog generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1.

Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 654–664.

A Supplemental Material

A.1 Training Details

Supervised Pre-train	
Word Embedding	256
Utterance Encoder	Attn GRU (128)
Context (De)Encoder	GRU (256)
Optimizer	Adam (lr=1e-3)
Dropout	0.5
β	0.01
Categorical \mathbf{z}	M=10, K=20
Gaussian \mathbf{z}	M=200
Reinforce	
Optimizer	SGD (lr=0.2 grad_clip=0.1)
γ	0.95

Table 7: Training details for DealOrNoDeal experiments. Attn GRU refers to (Yang et al., 2016)

Supervised Pre-train	
Word Embedding	256
Encoder	Attn GRU (300)
Decoder	LSTM (150)
Optimizer	Adam (1e-3)
Dropout	0.5
β	0.01
Categorical \mathbf{z}	M=10, K=20
Gaussian \mathbf{z}	M=200
Reinforce	
Optimizer	SGD(lr=0.01 grad_clip=0.5)
γ	0.99

Table 8: Training details for MultiWoz experiments

A.2 DealOrNoDeal Example Conversations

The followings are examples dialogs generated from word-level models (Table 9) and latent-level models (Table 10).

Scenario 1

1 book, 1 hat and 3 balls available.

Value Function

Agent: book=1, hat=6, ball=1.

User: book=1, hat=6, ball=1.

Dialog

Agent: I’d like the hat and 2 balls.

User: Deal.

Agent: SELECTION

Outcome: Agreement

Agent: 1 hat and 2 balls; User: 1 book and 1 ball

Agent reward = 8; User reward = 2

Scenario 2

4 books, 1 hat and 1 ball available.

Value Function

Agent: book=0, hat=7, ball=3;

User: book=1, hat=6, ball=0.

Dialog

User: I can offer you 1 hat and 2 books.

Agent: I need the hat and ball.

User: That won’t work for me. I can offer you one book and the hat.

Agent: I need the hat and ball.

User: I have to have the hat.

Agent: I need the hat and ball.

User: Okay.

Agent: SELECTION

Outcome: Agreement

Agent: 1 hat and 1 ball; User: 4 books

Agent reward = 10; User reward = 4

Table 9: Example dialogs between baseline with the user model. Agent is trained with word-level policy gradient and the user is a supervised pre-trained model.

Scenario 1

2 books, 2 hats and 1 ball available.

Value Function

Agent: book=3, hat=0, ball=4;

User: book=2, hat=3, ball=0.

Dialog

Agent: Can i have the books and one ball?

User: I would like the books and the hats then.

Agent: I would like the ball and books.

User: Deal.

Agent: SELECTION

Outcome: Agreement

Agent: 2 books and 1 ball

User: 2 hats

Agent reward = 10, User reward = 6

Scenario 2

2 books, 1 hat and 3 ball available.

Value Function

Agent: book=0, hat=4, ball=2;

User: book=3, hat=1, ball=1.

Dialog

Agent: I would like the hat and balls.

User: I cannot accept that. I need the hat and books.

Agent: Can I have the hat and balls, you get the rest?

User: Okay, deal.

Agent: SELECTION.

Outcome: Agreement

Agent: 1 hat and 3 ball

User: 2 books

Agent reward = 10, Simulator reward = 6

Table 10: Example dialogs between LiteCat and the user model. Agent is trained with latent-level policy gradient and the user is a supervised pre-trained model.