

Customized Nonlinear Bandits for Online Response Selection in Neural Conversation Models

Bing Liu,* Tong Yu,* Ian Lane, Ole J. Mengshoel

Electrical and Computer Engineering, Carnegie Mellon University
{liubing, lane}@cmu.edu, tongy1@andrew.cmu.edu, ole.mengshoel@sv.cmu.edu

Abstract

Dialog response selection is an important step towards natural response generation in conversational agents. Existing work on neural conversational models mainly focuses on offline supervised learning using a large set of context-response pairs. In this paper, we focus on online learning of response selection in retrieval-based dialog systems. We propose a contextual multi-armed bandit model with a nonlinear reward function that uses distributed representation of text for online response selection. A bidirectional LSTM is used to produce the distributed representations of dialog context and responses, which serve as the input to a contextual bandit. In learning the bandit, we propose a customized Thompson sampling method that is applied to a polynomial feature space in approximating the reward. Experimental results on the Ubuntu Dialogue Corpus demonstrate significant performance gains of the proposed method over conventional linear contextual bandits. Moreover, we report encouraging response selection performance of the proposed neural bandit model using the Recall@k metric for a small set of online training samples.

Introduction

Conversational agents, or chatbots, have a wide range of applications such as in technical support, personalized service, and entertainment (Young et al. 2013; Serban et al. 2016). Conventional approaches in conversational modeling include using rule-based and learning-based methods (Oh and Rudnicky 2000; Schatzmann et al. 2006) that typically require expert knowledge in designing rules and features. Such systems are usually designed to model dialogs in a specific domain. Transferring such models from one domain to another is often difficult.

With the explosive growth of conversational data in social media, neural network based data-driven approaches (Shang, Lu, and Li 2015; Vinyals and Le 2015) to conversational modeling have been explored. These models are trained based on large dialog corpora, making few assumptions about dialog domain or structures. They are suitable in modeling non-task-oriented conversations, and have shown promising performance in modeling social media chats (Ritter, Cherry, and Dolan 2011; Shang, Lu, and Li 2015; Li et

al. 2016b) and movie conversations (Vinyals and Le 2015; Li et al. 2016c). Most of the existing work focuses on the offline supervised dialog model training setting, in which the models are trained with a large number of context-response pairs. Online learning methods for neural conversational models, which are useful in building personalized dialog systems and transferring systems to new dialog domains, is not widely studied in literature.

In neural conversational models, response to a dialog context can either be generated directly by the model or can be selected from a repository of predefined responses. Generation-based models have advantages in producing responses that are more diverse. However, they often require a large number of training samples. Evaluating generation-based systems is also challenging, with recent studies (Liu et al. 2016) showing that many of the automatic evaluation methods for response generation are poorly correlated with human judgement. Retrieval-based models that select a dialog response from a candidate set based on a dialog context can be seen as an intermediate step towards generation-based systems. Response selection is similar to response generation when the set of candidate responses is very large (Lowe et al. 2015; 2016). There are well defined automatic evaluation metrics for response selection models.

In this work, we focus on online learning for retrieval-based conversational models. We propose a neural nonlinear bandit model (NNBM) that can be trained online for dialog response selection. Given a dialog context, the model selects a most likely response from a list of response candidates. Binary feedback is collected from users indicating whether or not the selected responses are appropriate. This binary feedback serves as the reward to the contextual bandits.

Our main contributions in this work are:

- We design a contextual bandit model, NNBM, that has a nonlinear reward function applied on selected dimensions of a polynomial feature space. The proposed method significantly outperforms contextual bandits with linear rewards that are widely used in online recommendations.
- We apply neural contextual multi-armed bandits to online learning of response selection in retrieval-based dialog models. To our best knowledge, this is the first attempt at combining neural network methods and contextual multi-armed bandits in this setting.

*These two authors contributed equally.

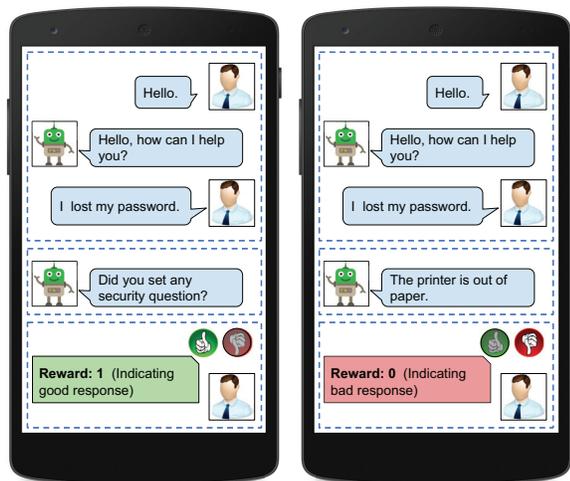


Figure 1: Illustration of the online model for response selection in dialog systems. Dialog history with a few number of turns is provided to the agent as the dialog context. The agent computes the most likely true response from a list of candidate responses and returns it to user. The user then provides binary feedback, positive (left) or negative (right), to the agent indicating the quality of the returned response.

Related Work

Conversation Modeling

Traditional approaches to conversation modeling typically involve template and rule based methods with statistical learning components (Levin and Pieraccini 1997; Oh and Rudnicky 2000). With the growth of conversational data in social media, data-driven approaches to conversation modeling have been widely studied. Ritter et al. (2011) framed conversation response generation as a statistical machine translation problem. Vinyals et al. (2015) designed a neural conversation model that can generate simple conversations and extract knowledge from an open-domain dataset. Shang et al. (2015) showed that the encoder-decoder-based neural network model could generate varied multiple responses to a given post on social networks. Serban et al. (2016) proposed a hierarchical recurrent neural model that utilizes context over an extended dialog history. Li et al. (2016a) proposed using maximum mutual information as the objective function to produce more diverse and interesting responses. These end-to-end neural conversation models are trained offline with supervised learning using large parallel corpora of utterance-response pairs. In this work, we focus on online learning of neural conversational models.

Reinforcement learning (RL) has also been applied in dialog modeling. In task-oriented dialog systems, RL can be applied to optimize dialog policy online with the feedback collected via interacting with users (Singh et al. 2000; Gašić et al. 2013). Zhao and Eskenazi (Zhao and Eskenazi 2016) introduced joint learning of dialog state tracking and dialog policy learning using deep RL. Li et al. (Li et al. 2017) proposed end-to-end trainable task-oriented dialog model with deep RL. Lipton et al. (2016) proposed

a Bayes-by-Backprop Q-network that effectively improves the efficiency of exploration with Thompson sampling. He et al. (2016) proposed a deep reinforcement relevance network (DRRN) that applies to text-based games in which action and state spaces are represented with separate embedding vectors and used to approximate the Q-function. These models are applied in a sequential decision problem setting for task-oriented dialogs. Our model, on the other hand, focuses on online learning in a setting where the context is provided and is not dependent on previous state and actions. The problem setting is closer to a non-task-oriented chit-chat dialog setting. In non-task-oriented dialog learning, Li et al. (2016c) introduced the use of RL for response generation by simulating dialogs between two agents that optimizes future reward. Serban et al. (2017) proposed a hybrid generation and retrieval based dialog model that applies RL in response selection. In these models, a simple reward function with linear approximation is applied in learning dialog policy. In this work, we investigate online learning efficiency using a more general nonlinear reward function approximation with Thompson sampling.

Evaluation of dialog system is difficult (Liu et al. 2016) as there is not yet a suitable metric for response generation evaluation. Many recent studies (Lowe et al. 2015; Kadlec, Schmid, and Kleindienst 2015) using the Ubuntu Dialogue Corpus instead proceed towards best response selection, which can be seen as an intermediate step towards response generation (Lowe et al. 2015) and has well accepted evaluation metrics. We use the same response selection methods and metrics in our study.

Bandits and Online Learning

Multi-armed bandit model is a popular online learning method with great success in many fields, such as online advertisement (Chapelle and Li 2011), recommender systems (Li et al. 2010), and stochastic combinatorial optimization (Gai, Krishnamachari, and Jain 2012; Kveton et al. 2014). The bandit enables online model update online based on the feedback from the environment, which has several nice properties in terms of finite-time regret bounds with a clear dependence on the parameters of interest. To model more complicated problems, a number of non-linear bandit models (Filippi et al. 2010; Gopalan, Mannor, and Mansour 2014; Kawale et al. 2015; Katariya et al. 2017b; 2017a; Yu, Kveton, and Mengshoel 2017) are proposed, in which the rewards are approximated by different non-linear functions. Bandits have also been applied in some dialog systems. Genevay et al. (2016) applied bandits in source user selection for user adaptation in spoken dialog systems. Bouneffouf et al. (2014) proposed active Thompson sampling method with bandits to select the most useful unlabeled examples to train a predictive model. They used a linear regression model to approximate the reward in each round of interaction. In our problem setting, the reward is a Bernoulli variable, where a logistic regression model is more suitable for binary reward approximation. Moreover, we apply bandits on distributed representation of the dialog context that are generated from neural network models, which is different from Bouneffouf et al. (2014).

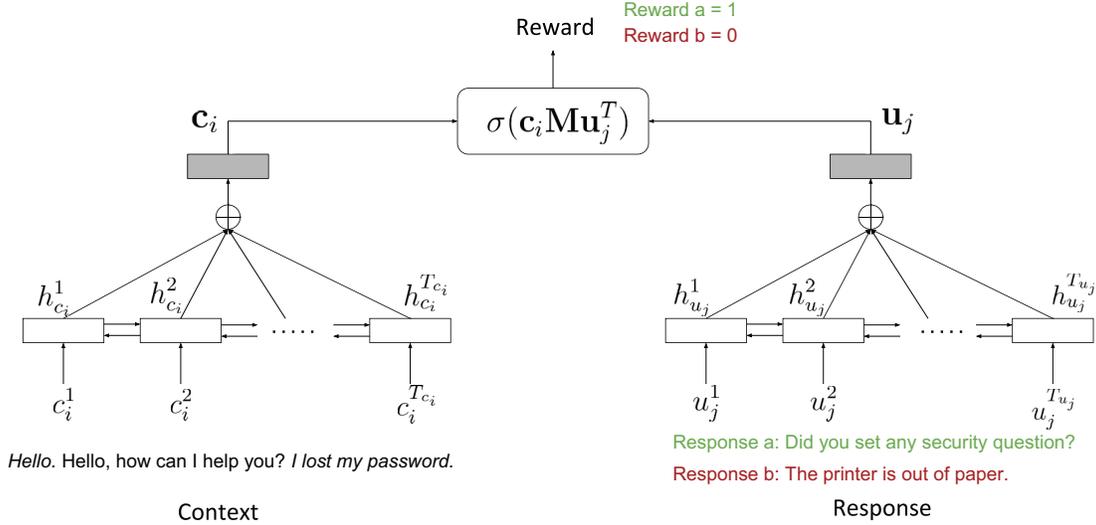


Figure 2: Architecture of the proposed Neural Nonlinear Bandit Model (NNBM). A bidirectional LSTM is used to encode dialog context and response to continuous representations. The context and response representations \mathbf{c}_i and \mathbf{u}_j are produced by taking the average of the LSTM state output sequence. \mathbf{c}_i and \mathbf{u}_j are then mapped to a second order polynomial feature space and serve as the input to the contextual bandits. A binary reward is finally collected from user indicating the quality of the selected response to the dialog context, which is used to update the NNBM parameters.

Neural Nonlinear Bandit Model

Figure 2 shows the architecture of our proposed neural nonlinear bandit model for online response selection in retrieval-based dialog systems. Let $C = \{c_1, c_2, c_3 \dots\}$ represent a set of dialog context samples and let $U = \{u_1, u_2, u_3 \dots\}$ represent a set of response samples in the dialog corpus.¹ A dialog context consists of a sequence of utterances by one or more users, and potentially by the NNBM. Given a dialog context c_i , the NNBM attempts to select a most likely response u_j from a pool of candidate responses. We ask the user to provide a binary feedback indicating whether or not the selected response is an appropriate one. This binary value r_{ij} serves as the reward to the contextual bandits during online learning.

We use a bidirectional LSTM to produce distributed representations of the dialog context and response. The context and response representations \mathbf{c}_i and \mathbf{u}_j are produced by taking the average of the LSTM sequence outputs $(h_{c_i}^1, \dots, h_{c_i}^{T_{c_i}})$ and $(h_{u_j}^1, \dots, h_{u_j}^{T_{u_j}})$. \mathbf{c}_i and \mathbf{u}_j serve as input features to the contextual bandit. In learning the bandit, we propose a Thompson sampling method that is applied to a second order polynomial feature space of the context and response in approximating the reward. We describe the proposed model details in the sections below.

Utterance Encoding with LSTM

We use a bidirectional LSTM to generate distributed feature representations for both the context and the response. For

¹In this paper, we use c_i to refer to the i th context sample and \mathbf{c}_i to denote the feature representation of the i th context sample. The same also applies to dialog responses.

the context reader, input to the network at each time step t is the embedding of the word c_i^t in context c_i . The corresponding output $h_{c_i}^t$ is the concatenation of the forward LSTM output $\vec{h}_{c_i}^t$ and backward LSTM output $\overleftarrow{h}_{c_i}^t$, i.e. $h_{c_i}^t = [\vec{h}_{c_i}^t, \overleftarrow{h}_{c_i}^t]$. The final context representation \mathbf{c}_i is produced by averaging over the sequence of LSTM state outputs. For the response reader, the same LSTM that shares the parameters with the context reader is used. The final response representation \mathbf{u}_j for response u_j is generated in the same manner as in the context representation generation.

The generated feature vectors \mathbf{c}_i and \mathbf{u}_j encode the meaning of the context and the response. During LSTM offline supervised pre-training, a score is generated for the context-response pair representing the probability of the match. The score of the context-response pair is calculated by applying the sigmoid function on the transformed inner product of \mathbf{c}_i and \mathbf{u}_j as:

$$p(\text{label}_{i,j} = 1 | \mathbf{c}_i, \mathbf{u}_j) = \sigma(\mathbf{c}_i \mathbf{M} \mathbf{u}_j^T) \quad (1)$$

where \mathbf{M} is a transformation matrix. In equation (1), the matrix \mathbf{M} is randomly initialized and updated during model training together with other neural network parameters. The neural network model is optimized by minimizing the cross entropy of all labeled context-response pairs in the training set:

$$\mathcal{L} = - \sum_{i,j} \log p(\text{label}_{i,j} = 1 | \mathbf{c}_i, \mathbf{u}_j) \quad (2)$$

Linear Logistic Regression Thompson Sampling

At round t of interaction in online model learning, a random context c_i is picked. Our agent chooses a user response u_j

and receives a binary reward (1 for like, and 0 for dislike). With context c_i and response u_j , our algorithm applies a parametric approximation of the reward r_{ij} by the representations c_i , u_j and unknown model parameters \mathbf{w} . As the reward in our problem setting is a Bernoulli variable, logistic regression is a suitable choice to model the reward.

In online advertising or recommendations with contextual bandits (Li et al. 2010; Chapelle and Li 2011), the reward r_{ij} is usually linear in feature representations of users, articles, or advertisements. In our case, we can also model r_{ij} by a linear logistic function with parameters \mathbf{w} , such that $r_{ij} = \sigma(\mathbf{x}_t \mathbf{w})$, where $\mathbf{x}_t = [c_i, u_j] \in \mathbb{R}^{1 \times 2L}$, $c_i \in \mathbb{R}^{1 \times L}$, $u_j \in \mathbb{R}^{1 \times L}$ and $\sigma(\cdot)$ is the sigmoid function.

Many algorithms have been studied in literature for balancing exploration and exploitation during policy learning, such as exploit-only, random, Thompson sampling (Chapelle and Li 2011), and upper confidence bound (UCB) (Auer, Cesa-Bianchi, and Fischer 2002). We apply Thompson sampling as it outperforms UCB, exploit-only, and random methods (Chapelle and Li 2011).

We first sample the parametrization of reward from its posterior. There is no closed-form solution for sampling from the posterior of logistic regression. Based on the Laplace’s method, the posterior at time t can be approximated (Spiegelhalter and Lauritzen 1990; MacKay 1992) by $\mathbf{w}_t \sim \mathcal{N}(\bar{\mathbf{w}}_{t-1}, S_t)$. The covariance matrix is $S_t = (X_t^T C_t X_t + \lambda I_L)^{-1}$, where C_t is a $t \times t$ diagonal matrix with diagonal elements $(C_t)_{ii} = \sigma(\mathbf{x}_i \bar{\mathbf{w}}_t)(1 - \sigma(\mathbf{x}_i \bar{\mathbf{w}}_t))$. Here $X_t = (\mathbf{x}_1^T, \dots, \mathbf{x}_t^T)^T \in \mathbb{R}^{t \times 2L}$ is the matrix of all context and response pair \mathbf{x}_i up time t , where $1 \leq i \leq t$.

To obtain a Gaussian approximation of the posterior distribution, we maximize the *logistic regression log-likelihood* $L_t(\mathbf{w})$ from the first t observations to get $\bar{\mathbf{w}}_t$, which is the mean of the Gaussian. We update $t(\mathbf{w})$ incrementally according to Newton’s method:

$$\bar{\mathbf{w}}_t = \bar{\mathbf{w}}_{t-1} - H_{t-1}^{-1}(\bar{\mathbf{w}}_{t-1}) \nabla(L_{t-1}(\bar{\mathbf{w}}_{t-1})),$$

where $\bar{\mathbf{w}}_{t-1}$ is the solution at time $t - 1$, $H_t(\mathbf{w})$ is the Hessian of the *logistic regression log-likelihood* $L_t(\mathbf{w})$ from the first t observations, and $\nabla(L_t(\mathbf{w}))$ is the corresponding gradient. The log-likelihood is defined as:

$$L_t(\mathbf{w}) = \lambda \mathbf{w}^T \mathbf{w} + \sum_{i=1}^t f_i(\mathbf{x}_i) \log(\sigma(\mathbf{x}_i \mathbf{w})) + \sum_{i=1}^t (1 - f_i(\mathbf{x}_i)) \log(1 - \sigma(\mathbf{x}_i \mathbf{w})),$$

where $f_i(\mathbf{x}_i)$ is the observed reward at time i and λ is the weight of the regularization term. After we sample \mathbf{w}_t , we choose the response with the highest expected reward with respect to \mathbf{w}_t :

$$\mathbf{u}_t = \arg \max_{u_t \in \text{all responses at } t} \sigma(\mathbf{x}_t \mathbf{w}_t),$$

where $\mathbf{x}_t = [c_t, u_t]$ and c_t is the context at t . Finally, we observe $f_t(\mathbf{x}_t)$ and receive it as a reward.

Logistic Regression Thompson Sampling on Polynomial Feature Space

We further propose a contextual bandit model with a nonlinear reward function applied on selected dimensions of a second order polynomial feature space. In approximating the reward, we apply the sigmoid function on the transformed inner product of \mathbf{c} and \mathbf{u} , $\mathbf{c} \mathbf{M} \mathbf{u}^T$:

$$[\mathbf{c}^{(1)} \quad \mathbf{c}^{(2)} \quad \dots \quad \mathbf{c}^{(L)}] \begin{bmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} & \dots & \mathbf{M}_{1L} \\ \mathbf{M}_{21} & \mathbf{M}_{22} & \dots & \mathbf{M}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{M}_{L1} & \mathbf{M}_{L2} & \dots & \mathbf{M}_{LL} \end{bmatrix} \begin{bmatrix} \mathbf{u}^{(1)} \\ \mathbf{u}^{(2)} \\ \vdots \\ \mathbf{u}^{(L)} \end{bmatrix} \quad (3)$$

where $\mathbf{c}^{(i)}$ and $\mathbf{u}^{(j)}$ are the i th and j th dimension of \mathbf{c} and \mathbf{u} . \mathbf{M}_{ij} indicates the element in the i th row and the j th column of \mathbf{M} , which is also the weight coefficient of term $\mathbf{c}^{(i)} \mathbf{u}^{(j)}$.

This parametric approximation is similar to a degree-2 polynomial kernel function in the space of \mathbf{x} , where $\mathbf{x} = [\mathbf{c}, \mathbf{u}] \in \mathbb{R}^{1 \times 2L}$. The only difference is that the weight coefficients of some terms are zeros. Let $\mathbf{x}^{(i)}$ be the i th dimension of \mathbf{x} . The coefficients of $(\mathbf{x}^{(i)})^2$ and $\mathbf{x}^{(i)}$ are zero where $1 \leq i \leq 2L$. Besides, the coefficients of $\mathbf{x}^{(j)} \mathbf{x}^{(k)}$ are zero where $1 \leq j \leq k \leq L$, or $L + 1 \leq j \leq k \leq 2L$. The bias term’s coefficient is zero as well. As a result, we approximate the reward by $\sigma(\phi(\mathbf{x}) \mathbf{w})$, where $\phi(\mathbf{x})$ maps \mathbf{x} into a higher-dimensional space. In addition, those zero coefficients guide us to remove some useless terms but select the important ones in the explicitly mapped degree-2 polynomial space. Therefore, this method can provide more powerful approximation compared to linear approximation studied previously (Li et al. 2010; Chapelle and Li 2011).

Experiments

We evaluate the efficiency of the proposed online dialog learning method on the average cumulative regret and response selection Recall@k metrics. We describe the experiment settings and evaluation metrics, and discuss the experimental results.

Data Set

We use the Ubuntu Dialogue Corpus (UDC) (Lowe et al. 2015) in our evaluation. UDC is a multi-turn based dialog corpus that is constructed from Ubuntu chat logs, used to receive technical support for Ubuntu-related problems. All named entities in the data set are replaced with corresponding tags (names, locations, organizations, URLs, and paths) in the data pre-processing stage. The data set contains dialog context-response pairs that are extracted and sampled from real chat logs. Each example contains three fields: the dialog context, the candidate response, and a label (match or non-match).

Evaluation Metrics

We evaluate our model and proposed learning methods on two different metrics, average cumulative regret and Re-

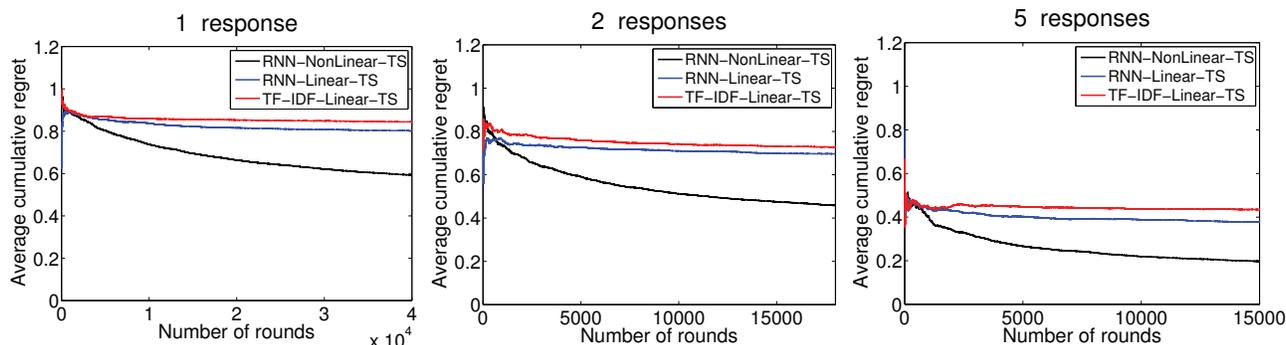


Figure 3: The average cumulative regret along the growing round of interactions in bandits for k ($k=1, 2, 5$) responses returned per interaction. k indicates the number of responses returned by the agent in each interactions. Corresponding rewards for the k responses are provided by the user. Total sample size for online learning is 1000.

	TF-IDF-Linear-TS	RNN-Linear-TS	RNN-NonLinear-TS
1 in 10 R@1	10.0 %	10.5 %	24.5 %
1 in 10 R@2	19.5 %	21.0 %	38.0 %
1 in 10 R@5	60.5 %	67.5 %	75.0 %

Table 1: Recall@ k evaluation results for the proposed bandit models. TF-IDF-Linear-TS, RNN-Linear-TS and RNN-NonLinear-TS are evaluated when $k = 1, 2$ and 5 .

call@ k . The average cumulative regret is defined as:

$$R(T) = \frac{\sum_{t=1}^T r_{t,u_t^*} - \sum_{t=1}^T r_{t,u_t}}{T}, \quad (4)$$

where r_{t,u_t^*} is the optimal reward could be achieved at round t , and r_{t,u_t} is the reward achieved by bandits at round t .

For the evaluation using Recall@ k , the agent is provided with 10 candidate responses. The goal of the agent is to rank the responses by assigning a high score to the true response, and assigning lower scores to the false or distracting responses. Recall@ k means that the ranking is considered correct if the true response is among the top k selections made by the agent out of the 10 candidate responses. We report Recall@ k with $(n, k) \in \{(10, 1), (10, 2), (10, 5)\}$.

Experiment Settings

We use a bidirectional LSTM to encode dialog context and response as it has stronger capability in capturing longer term temporal dependencies (Hochreiter and Schmidhuber 1997) comparing to vanilla RNN. LSTM state size and output size are both set as 128. Word embeddings of size 150 are randomly initialized and fine-tuned during mini-batch (size 128) training. We use Adam optimizer (Kingma and Ba 2014) in the neural network offline model training with initial learning rate of $1e-3$. Dropout (Srivastava et al. 2014) with keep probability of 0.5 is applied during offline supervised pre-training.

We compare three different bandit models in the main experimental results:

- TF-IDF-Linear-TS: The context and response features are generated by TF-IDF followed by a dimension reduction operation with PCA. A linear logistic func-

tion $\sigma(\mathbf{x}\mathbf{w})$ is used to approximate the reward, where $\mathbf{x} = [\mathbf{c}, \mathbf{u}]$.

- RNN-Linear-TS: The dialog context and response are encoded with a bidirectional LSTM. As in TF-IDF-Linear-TS, a linear logistic function $\sigma(\mathbf{x}\mathbf{w})$ is used to approximate the reward.
- RNN-NonLinear-TS: This is our proposed NNBM. The dialog context and response are encoded with a bidirectional LSTM. A nonlinear logistic function $\sigma(\mathbf{c}\mathbf{M}\mathbf{u}^T)$, which is also $\sigma(\phi(\mathbf{x})\mathbf{w})$, is used to approximate the reward.

We further compare our method to Thompson sampling with linear regression proposed in (Bouneffouf et al. 2014), as it is the most closely related recent work using bandits with Thompson sampling for dialog learning. Bouneffouf et al. use a linear bandit with Thompson sampling to handle continuous reward. Our method uses non-linear bandit with Thompson sampling to handle Bernoulli reward. In comparing these two bandit methods, we apply the same text encoding using a bidirectional LSTM.

In model training, the bidirectional LSTM encoder is trained using data from the original UDC training set. Online bandit learning and evaluation is performed using data sampled from the UDC test set.

Online Learning Evaluations

We evaluate the performance of different bandit models in the online setting. We select 1000 samples from the UDC test set for the online model evaluation. In each round of interaction, a context is randomly selected from the 1000 samples and provided to the agent. Given the context, the agent returns k responses to the user. The user then provides

Dialog Context:

User A:	what you looking for linuxuz3r? no i mean are you looking for a specfic program? im not sure if there is anything better then sourceforge
User B:	no particular program, anything that interest me then contribute to the source. i wanna learn how to read code

Interaction 1:

Rank	Confidence	Response
1	0.58374	other then checking the addational-drivers tool (jockey-gtk) and see if you got the drivers installed.. thats all i know about ati cards.
2	0.46708	we needed some custom features not all the features of ffmpeg so i needed those features only.
3	0.37162	good lad

Interaction 4:

Rank	Confidence	Response
1	0.62103	there is one that escapes me at the moment. most people use sourceforge
2	0.46708	possible
3	0.46526	we needed some custom features not all the features of ffmpeg so i needed those features only.

Interaction 7:

Rank	Confidence	Response
1	0.87181	there is one that escapes me at the moment. most people use sourceforge
2	0.17525	giles that is the exact command, replace username with the actual name of the account you are trying to change. ex. 'sudo passwd kriskropd'
3	0.10332	we needed some custom features not all the features of ffmpeg so i needed those features only.

Table 2: An example showing how bandit model learns with increasing number of interactions. In each round of interactive learning, a context is randomly selected as input. Here we show the response ranking to one particular context sample. Interaction 1, 4, 7 indicates the round of interactions when this particular context appears for the 1st, 4th and 7th time respectively. For each response, logistic regression model in bandits provides a confidence score, based on which we sort all the response and show the top three. The response in bold represents the true response.

a binary (1 or 0) feedback on each of the returned response, where 1 indicates a good response, and 0 indicates a bad response. If a bad response is selected, the regret value, which is initialized to 0, increments by 1. We report the average cumulative regret in Equation 4 over 50000 rounds of interactions.

RNN-Linear-TS versus TF-IDF-Linear-TS. The results are reported in Figure 3. As can be observed from these figures, with increasing number of rounds of interactions in bandit model learning, the average cumulative regret gradually decreases and finally converges. Comparing to RNN based text encoding method, TF-IDF as a count-based method that runs faster and requires much less computation. However, it ignores the word ordering and does not learn word relations in the sentence. For bandits using linear logistic function, RNN-Linear-TS outperforms TF-IDF-Linear-TS in all experiments. This shows the advantages of using RNN methods for text feature representation that captures long term temporal dependencies in text comparing to conventional methods using TF-IDF.

RNN-NonLinear-TS versus RNN-Linear-TS. As illustrated in Figure 3, RNN-NonLinear-TS beats RNN-Linear-TS by large margin at convergence in

all experiment settings, at a cost of higher computational complexity. The proposed nonlinear logistic regression Thompson sampling method produces much better results in this online dialog response selection task, compared to the linear models that are widely applied in online advertisements. With total sample size of 1000 (Figure 3) and a single returned response (i.e. $k = 1$), the average cumulative regret converges at 0.61, indicating the probability of bandits selecting correct response is $0.39 (= 1 - 0.61)$ in each round.

We further compare our RNN-NonLinear-TS methods to previous work on using linear regression (Bounefouf et al. 2014) for reward approximation in dialog systems. In comparing these two bandit methods, we used the same text encoding by LSTM. As illustrated in Figure 4, our proposed NNBM method demonstrates strong advantage in online learning efficiency comparing to previously proposed method.

Table 2 shows an example showing how bandit model learns with increasing number of online learning interactions. It can be seen that with more rounds of interactions, the model confidence on the true response increases and the gap between the top two confidence scores expands.

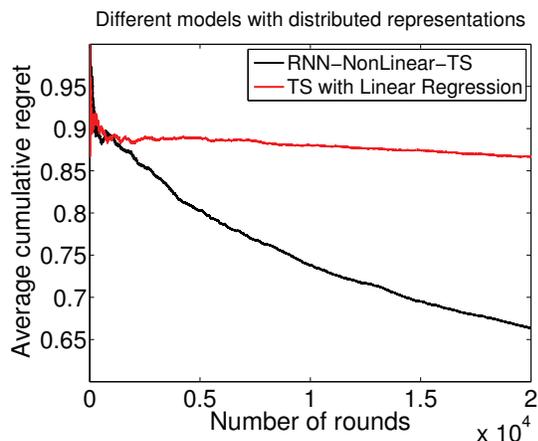


Figure 4: Comparison between RNN-NonLinear-TS and Thompson sampling with linear regression (Bouneffouf et al. 2014) in approximating the reward.

On number of responses. We experiment with letting the agent to select different number of responses for each given dialog context, with k value being 1, 2 and 5. The results are shown in Figure 3. As expected, the true response is more likely to appear in the returned responses when the number of candidate responses k increases, and thus resulting in a lower average cumulative regret. The gap between RNN-NonLinear-TS and RNN-Linear-TS expands when k grows from 1 to 2, and shrinks when k grows further to 5. This might be explained by the power of RNN-NonLinear-TS model in ranking the true response higher among the candidate responses.

Recall@k Results

We further evaluate the performance of our dialog response selection model using Recall@k. We split our 1000 learning samples into two parts: 800 samples are used in the online learning by bandits, and the rest 200 samples are set aside for evaluation. The results are shown in Table 1. Note that these numbers should not be directly compared to the ones in (Lowe et al. 2015) and (Kadlec, Schmid, and Kleindienst 2015), which are reported under offline model training settings using the entire training set. These Recall@k evaluation results are promising given the very small size (800) of samples involved in the online model learning. As in the above online evaluation results, RNN-Linear-TS outperforms TF-IDF-Linear-TS for all k values under the Recall@k metric. RNN-NonLinear-TS consistently outperforms RNN-Linear-TS by a large margin.

Conclusions

In this paper, we proposed a neural nonlinear bandit model (NNBM) using distributed representations of text in contextual multi-armed bandits for online response selection in dialog modeling. We designed a customized Thompson sampling method that is applied to a polynomial feature space in approximating the reward for bandits. Our experiments

results on Ubuntu Dialogue Corpus showed clear advantages of using distributed representation of text produced by neural network methods in learning text representations by comparing to using TF-IDF method for online dialog model learning. The experimental results demonstrated significant performance gain of the proposed methods over conventional linear contextual bandits, which are widely used in online advertisement and recommendations. We also reported encouraging Recall@k evaluation results of NNBM in best response selection.

To our best knowledge, this is the first attempt at combining neural network methods and contextual multi-armed bandits in the dialog response selection online learning setting. We believe the proposed online learning method can be very useful in situations where the size of labeled data set is limited.

Acknowledgments

This work is partially supported by Adobe donation. We thank Branislav Kveton for the helpful discussions.

References

- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47:235–256.
- Bouneffouf, D.; Laroche, R.; Urvoy, T.; Féraud, R.; and Allesiardo, R. 2014. Contextual bandit for active learning: Active thompson sampling. In *International Conference on Neural Information Processing*, 405–412. Springer.
- Chapelle, O., and Li, L. 2011. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, 2249–2257.
- Filippi, S.; Cappe, O.; Garivier, A.; and Szepesvári, C. 2010. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, 586–594.
- Gai, Y.; Krishnamachari, B.; and Jain, R. 2012. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking* 20(5):1466–1478.
- Gašić, M.; Breslin, C.; Henderson, M.; Kim, D.; Szummer, M.; Thomson, B.; Tsiakoulis, P.; and Young, S. 2013. Online policy optimisation of bayesian spoken dialogue systems via human interaction. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 8367–8371. IEEE.
- Genevay, A., and Laroche, R. 2016. Transfer learning for user adaptation in spoken dialogue systems. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 975–983. International Foundation for Autonomous Agents and Multiagent Systems.
- Gopalan, A.; Mannor, S.; and Mansour, Y. 2014. Thompson sampling for complex online problems. In *International Conference on Machine Learning*, 100–108.
- He, J.; Chen, J.; He, X.; Gao, J.; Li, L.; Deng, L.; and Ostendorf, M. 2016. Deep reinforcement learning with a natural language action space. In *ACL*.

- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Kadlec, R.; Schmid, M.; and Kleindienst, J. 2015. Improved deep learning baselines for ubuntu corpus dialogs. *arXiv preprint arXiv:1510.03753*.
- Katariya, S.; Kveton, B.; Szepesvári, C.; Vernade, C.; and Wen, Z. 2017a. Bernoulli rank-1 bandits for click feedback. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2001–2007.
- Katariya, S.; Kveton, B.; Szepesvári, C.; Vernade, C.; and Wen, Z. 2017b. Stochastic rank-1 bandits. In *Artificial Intelligence and Statistics*, 392–401.
- Kawale, J.; Bui, H. H.; Kveton, B.; Tran-Thanh, L.; and Chawla, S. 2015. Efficient thompson sampling for online matrix-factorization recommendation. In *Advances in Neural Information Processing Systems*, 1297–1305.
- Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. In *ICLR*.
- Kveton, B.; Wen, Z.; Ashkan, A.; Eydgahi, H.; and Eriksson, B. 2014. Matroid bandits: Fast combinatorial optimization with learning. In *UAI-2014*, 420–429.
- Levin, E., and Pieraccini, R. 1997. A stochastic model of computer-human interaction for learning dialogue strategies. In *Eurospeech*, volume 97, 1883–1886.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, 661–670. ACM.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL-HLT*.
- Li, J.; Galley, M.; Brockett, C.; Spithourakis, G. P.; Gao, J.; and Dolan, B. 2016b. A persona-based neural conversation model. In *ACL*.
- Li, J.; Monroe, W.; Ritter, A.; and Jurafsky, D. 2016c. Deep reinforcement learning for dialogue generation. In *EMNLP*.
- Li, X.; Chen, Y.-N.; Li, L.; and Gao, J. 2017. End-to-end task-completion neural dialogue systems. *arXiv preprint arXiv:1703.01008*.
- Lipton, Z. C.; Gao, J.; Li, L.; Li, X.; Ahmed, F.; and Deng, L. 2016. Efficient exploration for dialogue policy learning with bbq networks & replay buffer spiking. *arXiv preprint arXiv:1608.05081*.
- Liu, C.-W.; Lowe, R.; Serban, I. V.; 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 *EMNLP*.
- Lowe, R.; Pow, N.; Serban, I.; and Pineau, J. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. *arXiv preprint arXiv:1506.08909*.
- Lowe, R.; Serban, I. V.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. On the evaluation of dialogue systems with next utterance classification. In *SIGDIAL*.
- MacKay, D. J. 1992. The evidence framework applied to classification networks. *Neural computation* 4(5):720–736.
- Oh, A. H., and Rudnicky, A. I. 2000. Stochastic language generation for spoken dialogue systems. In *Proceedings of the 2000 ANLP/NAACL Workshop on Conversational systems-Volume 3*, 27–32. Association for Computational Linguistics.
- Ritter, A.; Cherry, C.; and Dolan, W. B. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, 583–593. Association for Computational Linguistics.
- Schatzmann, J.; Weilhammer, K.; Stuttle, M.; and Young, S. 2006. A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *The knowledge engineering review* 21(02):97–126.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; and Pineau, J. 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)*.
- Serban, I. V.; Sankar, C.; Germain, M.; Zhang, S.; Lin, Z.; Subramanian, S.; Kim, T.; Pieper, M.; Chandar, S.; Ke, N. R.; et al. 2017. A deep reinforcement learning chatbot. *arXiv preprint arXiv:1709.02349*.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. In *ACL-IJCNLP*.
- Singh, S.; Kearns, M.; Litman, D. J.; Walker, M. A.; et al. 2000. Empirical evaluation of a reinforcement learning spoken dialogue system. In *AAAI/IAAI*, 645–651.
- Spiegelhalter, D. J., and Lauritzen, S. L. 1990. Sequential updating of conditional probabilities on directed graphical structures. *Networks* 20(5):579–605.
- Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1):1929–1958.
- Vinyals, O., and Le, Q. 2015. A neural conversational model. In *ICML*.
- 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.
- Yu, T.; Kveton, B.; and Mengshoel, O. J. 2017. Thompson sampling for optimizing stochastic local search. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer.
- Zhao, T., and Eskenazi, M. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *SIGDIAL*.