
Learning Robust Dialog Policies in Noisy Environments

Maryam Fazel-Zarandi
Amazon Inc, USA
fazelzar@amazon.com

Shang-Wen Li
Amazon Inc, USA
shangwel@amazon.com

Jin Cao
Amazon Inc, USA
jincao@amazon.com

Jared Casale
Amazon Inc, USA
casjared@amazon.com

Peter Henderson
Amazon Inc, USA
pehender@amazon.com

David Whitney
Brown University
dwhitney@cs.brown.edu

Alborz Geramifard
Amazon Inc, USA
alborzg@amazon.com

Abstract

Modern virtual personal assistants provide a convenient interface for completing daily tasks via voice commands. An important consideration for these assistants is the ability to recover from automatic speech recognition (ASR) and natural language understanding (NLU) errors. In this paper, we focus on learning robust dialog policies to recover from these errors. To this end, we develop a user simulator which interacts with the assistant through voice commands in realistic scenarios with noisy audio, and use it to learn dialog policies through deep reinforcement learning. We show that dialogs generated by our simulator are indistinguishable from human generated dialogs, as determined by human evaluators. Furthermore, preliminary experimental results show that the learned policies in noisy environments achieve the same execution success rate with fewer dialog turns compared to fixed rule-based policies.

1 Introduction

Modern speech-based assistants, such as Amazon Alexa, Apple Siri, and Google Assistant, enable users to complete daily tasks such as shopping, setting reminders, and getting answers to factual questions using vocal commands. Such human-like interfaces create a rich experience for users by enabling them to complete many tasks hands-free and eyes-free in a conversational manner. Depending on how complicated a task is, multiple rounds of conversation may be needed for the assistant to fully understand user requests. To achieve this, these interfaces require the design of complex dialog policies which can generate appropriate responses to queries and steer the conversation.

Designing dialog policies in current speech-based assistants is challenging and time-consuming. In particular, responding to a user can be difficult since the selection of system actions is conditioned on the task, the user goals and preferences, and the dialog history. Additionally, environment noise and ambiguous user utterances present added complexity. Communicating with speech-based assistants always involves a noisy communication channel and the assistant should be able to recover from automatic speech recognition (ASR) and natural language understanding (NLU) errors by confirming its recognition results or eliciting more information from the user. To address these issues, both rule-based strategies and statistical modeling techniques have been developed [27]. Rule-based

systems are costly to design and maintain, and generally assume the dialog state is fully observable with only a limited account of errors and uncertainties. Statistical modeling techniques, on the other hand, are able to generalize to such unknowns, but rely on the availability of huge amounts of training data. Some statistical-modeling techniques formulate dialog management as a reinforcement learning (RL) problem [16]. In this formulation, the goal is to adjust a parameterized dialog policy in order to maximize the expected cumulative reward over the course of a dialog. To reduce the cost of collecting the large amount of data needed by such models, researchers have focused on the development of user simulators which interact with the dialog policy (e.g., [6, 8, 9, 12, 27, 29]). Most of the existing simulators are domain dependent [9] and lack robustness against noise in real user utterances in order to recover from ASR/NLU errors [19].

In this paper, we focus on learning dialog policies in noisy environments where the assistant needs to recover from ASR/NLU errors. The key contribution of this paper is the development of a realistic user simulator which interacts with the assistant through voice commands. The simulator is trained on real-world data and takes dialog context into account. Furthermore, while other works simulate NLU errors at the intent- or slot-level [18], our simulator is able to use artificial noise in audio signals to induce a similar corpus-specific ASR word-error rate (WER). This error can then propagate to realistic intent- or slot-level errors. Using dialogs collected from MovieBot [25], an Alexa Skill [14] that converses with users about existing, new, and forthcoming movies, we show that our simulator generates dialogs that are indistinguishable from human generated dialogs. Using the simulator, we leverage deep RL to learn robust dialog policies that overcome upstream ASR/NLU errors in a noisy environment. Experimental results show that the learned policies achieve the same execution success rate with fewer dialog turns compared to fixed rule-based policies in noisy environments.

2 Related Work

A dialog system can be formalized as a Markov Decision Process (MDP) [16]. An MDP is a tuple $\langle S, A, P, R, \gamma \rangle$ of states, actions, transition probability function, reward, and discount factor. In the formulation of RL, at each time step t , the agent observes a state $s_t \in S$ and selects an action $a_t \in A$ according to its policy ($\pi : S \rightarrow A$). After performing the selected action, the agent receives the next state s_{t+1} and a scalar reward r_t . The trajectory restarts after the agent reaches a terminal state. With this formalization, RL can be used to find the optimal dialog policy (e.g., [26, 34, 15, 13]). In this context, at each turn the system acts based on its understanding of what the user has said, and the reward function is modeled in terms of different dimensions such as per-interaction user satisfaction, accomplishment of the task, efficiency of interaction, dialog duration, etc.

Recently, deep RL has also been applied to the problem of dialog management. Cuayahuitl [7] presented an implementation of DQN [23, 24]. Zhao and Eskenazi [36] similarly used a deep recurrent Q-network (DRQN) to learn both state tracking and dialog policy for task-oriented dialogs. Fatemi et al. [10] compared different deep RL techniques on a restaurant-finding dialog problem. They showed that DDQN [32] converges faster than DQN for this domain.

An important challenge in using RL for learning dialog policies is creating realistic user simulators that can generate natural conversations similar to a human user [27]. Existing approaches for simulating users can be categorized into two groups: intent-level and utterance-level simulation. An intent-level simulator abstracts a conversation as a sequence alternating between user intents and slots and bot actions, where an intent indicates the user’s intention (e.g., getting the plot) and slots represent information about a particular entity (e.g., movie title). The goal is to predict the next user intent based on the dialog history (e.g., [8, 11, 18, 28]). These methods require extensive work to manually design the user goal or agenda, which makes them less generalizable. Additionally, it is difficult to simulate real-world recognition errors. Utterance-level user simulators, on the other hand, attempt to generate natural language user responses. To this end, template-based natural language generation is usually adopted to output utterances based on handcrafted rules and predicted user intents [18]. Inspired by recent studies in sequence-to-sequence machine translation [31], end-to-end dialog simulators generate utterances with minimal feature engineering [9, 5]. However, these techniques often produce generic responses [17]. Additionally, they have to be trained with a significant amount of dialogs, and thus they do not work well for domain-specific applications.

Given that the dialog state is not fully observable due to ASR/NLU errors and uncertainty over a user’s goals, researchers have also investigated the use of Partially Observable Markov Decision Processes

(POMDPs) [35]. However, these systems are complex and exact policy learning for POMDPs is intractable [35]. Rule-based policies, on the other hand, use ASR/NLU confidence scores to recover from errors. However, asking too many clarifying questions can frustrate users and in some scenarios it might be more beneficial to execute on available information. Additionally, ASR/NLU components are not static and as such policies with fixed confidence levels need to be constantly updated.

3 User Simulation

To develop a user simulator that generates voice responses to the agent, we adopt a hybrid of language model based intent prediction and template-based utterance generation. The overall architecture of the user simulator is illustrated in Figure 1.

The user simulator first generates the next intent and slot type based on dialog context. Dialog context includes previous intent and slot type pairs in the conversation, as well as text responses from the agent encoded with word2vec [22]. We model intent generation as a language modeling problem and use recurrent neural networks (RNNs) to predict the trajectory of intents. In this formulation, each possible intent and slot type pair form a token in the vocabulary, and every training dialog becomes a training intent sequence. For example, the sequence for the conversation with MovieBot in Table 1 is *[Start, GetGenreMoviesIntent+genre, GetNextActorIntent, GetRatingIntent, End]*.

Given the predicted intent and slot type, the user simulator uniformly samples one utterance from the templates and fills the utterance with slot values uniformly sampled from a lexicon of the corresponding slot type. For example, given intent *GetDirectorIntent* and slot type *movie_title*, a sample utterance from the template is "who is the director of {movie_title}?". {movie_title} is then replaced uniformly by a value from the lexicon of movie titles. When the simulator is responding to a clarifying question, it answers yes or no to a confirmation question, but repeats the original utterance for elicitation.

Providing the sampled text directly to the MovieBot NLU model would result in practically zero errors as the sample utterances are used in training the NLU model via the ASK portal. Therefore, we take the approach of generating a spoken audio signal from the text, injecting noise into the audio and then passing the audio to ASR and NLU models of the MovieBot Skill. This process simulates the real-world sources of error such as background noise in a person’s home. Concretely, we synthesize an audio signal with the text-to-speech (TTS) service provided by Amazon Polly [1] and contaminate the clean TTS output using PyAcoustics [20] to achieve a desired signal-to-noise ratio (SNR). Contamination is performed by filtering a white noise signal using a Fast Fourier transform applied to the original TTS output. Given our choice of architecture, any noise injection strategy may be developed without impacting the rest of the experimental setup. For example, in the future we plan

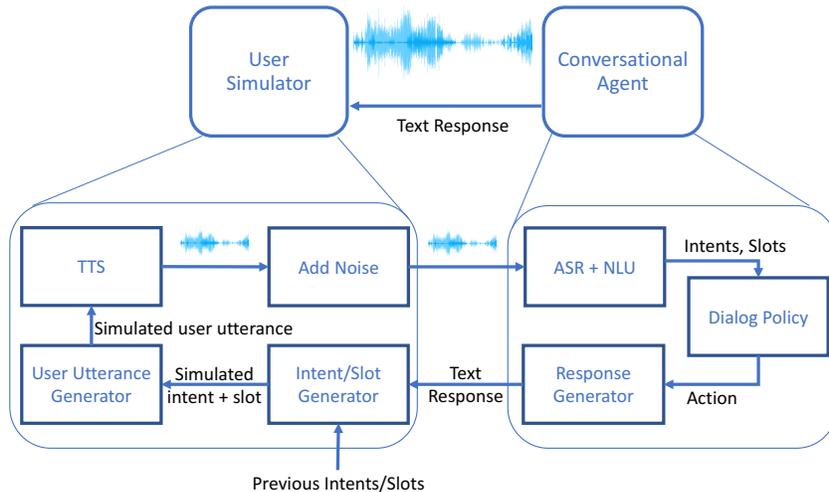


Figure 1: User simulator and conversational agent interaction. We use the text response generated by the agent before sending it to TTS to eliminate the need for ASR/NLU on the user simulator side.

Table 1: Example of a conversation with MovieBot. The intent indicates the user’s intention and slots (e.g., genre) represent information about a particular entity.

Speaker	Utterance
User	A science fiction movie [intent = GetGenreMoviesIntent, genre = science fiction]
Bot	How about Arrival, starring Amy Adams. If you want, you can say who else is in it.
User	Who else is in it? [intent = GetNextActorIntent]
Bot	The character of Ian Donnelly was played by Jeremy Renner.
User	Is it any good? [intent = GetRatingIntent]
Bot	It is rated 8.4 on I.M.D.b., based on about 70000 votes.
User	Thank you [intent = AMAZON.StopIntent]
Bot	Thank you for using MovieBot.

to explore the impact of using different kinds of noise such as background noise, speaker-dependent speech variations, or multiple concurrent speakers. In this particular set of experiments, we tune the injected white noise to a level that causes the ASR WER of the simulator to match that of MovieBot’s deployed ASR model.

This approach to simulating users has a number of benefits. By predicting intents first, the model becomes highly accurate and data efficient at the intent level due to the smaller number of intents compared to words in the language. By using RNNs, our simulator learns long-term dependencies between dialog context and intents from data to predict consistent user behavior. Additionally, using handcrafted templates ensures that the generated utterances are of high quality, while the TTS component with synthesized noise allows us to simulate real-world noisy environments.

3.1 Implementation and Evaluation

We used a corpus of 966 MovieBot dialogs consisting of 6349 user turns to train the intent generation model with and without agent responses. The corpus is collected from real user interactions with MovieBot and annotated at the utterance level. For example, a user request such as "How old is Tom Hanks?" is annotated as intent=*GetAgeIntent* and slot *PersonName=Tom Hanks*. For intents and slot types, we added an embedding layer initialized with small random values. We used a single layer network followed by softmax, and for the hidden units we used two variants: vanilla RNN and Gated Recurrent Unit (GRU) [4]. The output of the network is a sequence of vectors, each with total number of intent/slot pair elements. The optimal parameters were found using grid search (Appendix A). For all models we set the learning rate to 0.001, and used Google’s trained word2vec model for encoding agent responses. As a baseline, we used a bigram model. Note that this model only takes the intent and slot type inputs. With this setting, we found that the vanilla RNN models have the lowest perplexity [30] scores across various trials (Table 2).

We also conducted a user study to evaluate if participants could distinguish simulator generated dialogs from real human dialogs. The participants were each presented 10 dialogs via a web interface and asked to determine if each dialog was generated by a human or a simulated user. Optionally, the participants could also provide feedback as to why they made their decision. For each participant, five dialogs were sampled from the annotated MovieBot data, and the first utterance of each was used to generate five dialogs with the simulator. We used two models for the simulated data, namely, the bigram model and the vanilla RNN model (see Appendix B for example of simulated dialogs). The difference between the true positive rate (human dialog correctly identified as human) and false positive rate (simulated user incorrectly identified as human) is of particular interest: more individuals identifying simulated user dialogs as human indicates better user simulation.

Table 2: Perplexity results for the Bigram, RNN, and GRU models. Note that a lower value indicates better performance.

Model	Bigram	RNN	GRU
Intent and Slot Only	16.270	7.675	7.769
Intent and Slot + Bot Response	–	7.274	7.892

Table 3: Results of the user study for distinguishing between real and simulated dialogs. The difference between identifying human as human and simulator as human is small.

	All Participants		Participants Who Left Feedback	
Number of Subjects	31		13	
Total Number of Dialogs	309		129	
Accuracy (Human/RNN + Bigram Models)	50.49%		58.14%	
Precision (Human/RNN + Bigram Models)	50.29%		57.14%	
Recall (Human/RNN + Bigram Models)	57.14%		62.50%	
F1-Score (Human/RNN + Bigram Models)	53.50%		59.70%	
Human Identified as Human	57.14%	154 dialogs	62.50%	64 dialogs
Simulator Identified as Human (RNN + Bigram Models)	56.13%	155 dialogs	46.15%	65 dialogs
Simulator Identified as Human (RNN Model)	59.21%	76 dialogs	58.82%	34 dialogs
Simulator Identified as Human (Bigram Model)	53.16%	79 dialogs	32.26%	31 dialogs

In total, 31 individuals from the Amazon Alexa team participated in the study and evaluated 309 dialogs. The results are reported in Table 3. Overall results indicate that the difference between identifying human as human (57.14%) and simulator as human (56.13%) is not statistically significant (p -value = 0.858). Additionally, the RNN model is performing slightly better than the bigram model. Upon further inspection, we found a weak correlation (Pearson’s correlation coefficient: 0.38, p -value: 0.036) between how participants rated a dialog and whether or not they left feedback. The results for the 13 participants who left feedback is reported separately in the table, clearly showing the RNN (58.82% of simulated users identified as human) significantly (p -value = 0.033) outperforms the bigram model (32.26% of simulated users identified as human). We hypothesize that this is due to bigram dialogs being longer on average, and the participants who left feedback read the dialogs more carefully. Criteria for selecting a dialog as a real human as indicated in the participant feedback included: the expression of emotions (e.g., the user gets mad when the bot doesn’t answer correctly), logical progression of the dialog, specific goals that humans have regardless of the dialog system’s response, and naturalness of the dialogs. On the other hand, top criteria for identifying simulated users included: asking a question without context, referring to non-existent entities in the dialog, random responses by the user, and the user not getting frustrated. These comments illustrate that better context tracking and use of sentiments would further improve user simulation. Nevertheless, given the small difference between true positive and false positive results, the user simulator can sufficiently replace real users for dialog policy learning.

4 Dialog Policy Learning

We used the proposed user simulator to train dialog policies through deep RL with the goal of conversational error recovery – that is, recovering from ASR/NLU errors through necessary clarifications. The components of the dialog policy learning problem are as follows:

- *S*: The state is composed of 1) the hypothesis intent and slot, 2) the ASR and NLU intent and slot confidence scores, and 3) the previous bot action.
- *A*: We constrain the action space to three actions that are critical for conversational error recovery: *execute*, *confirm*, and *elicit*. *execute* accesses an IMDb database and answers the user’s question. *confirm* and *elicit* are used to recover from ASR/NLU errors. *confirm* clarifies the intent and/or the slot with the user (e.g., "Do you want movies directed by Christopher Nolan?"), whereas *elicit* asks the user for missing information (e.g., "Which movie are you talking about?").
- *R*: Table 4 summarizes the reward function. Here, *ref* indicates the reference intent and slot as generated by the user simulator, and *hyp* refers to what the bot understands. If $hyp \neq ref$, this indicates that an NLU error was induced by added noise to the audio signal. The environment gives a large positive reward for executing correctly, a large negative reward for executing incorrectly, and smaller negative rewards for confirm and elicit. A successful execution is when the reference and hypothesis intent and slots are equal. Clarification, although sometimes necessary, can frustrate users, hence the small

Table 4: Reward Function

a	s	r
execute(hyp)	$hyp = ref$	+10
execute(hyp)	$hyp \neq ref$	-12
confirm	*	-3
elicit	*	-6
terminate (after correct execute)	*	+1
terminate (after incorrect execute)	*	-5

negative rewards for clarifying actions. Additionally, if the user leaves the conversation after a successful execution, the bot receives a small positive reward, but if the user leaves after an incorrect execution, the bot receives a medium negative reward. The reason we give a small positive reward to termination after a successful execution is that in the absence of an explicit goal we assume the user received the requested information and decided to terminate the dialog. All specific values were assigned through empirical analysis of user interactions with the simulator.

For the deep RL experiments we used DQN and Dueling DDQN [33], with a fully-connected Multi-Layer Perceptron (MLP) to represent the deep Q-network. The input of the network is the concatenated embedding vector of intent and slot, the vector of confidence scores, and previous bot action. Additionally, we tuned a window size to include previous dialog turns as input. The hidden layers use a rectifier nonlinearity, and the output layer is a fully connected layer with linear activation function and a single output for each valid action. We trained the agents using an ϵ -greedy policy with ϵ decreasing linearly from 1 to 0.1 over 100,000 steps. We ran each method 30 times for 150,000 steps, and in each run, after every 10,000 steps, we sampled 30 dialog episodes with no exploration to evaluate performance. All methods used the same set of random seeds and the best parameters were empirically found for each method using Hyperopt [2]. Additionally, we experimented with two fixed policies: 1) *execute* only and 2) *execute* with *confirm* and *elicit*. For the second fixed policy, we learned global ASR and NLU confidence score thresholds for MovieBot by optimizing the simulated agent’s average return (Table 5).

Table 5: Fixed Policy (MovieBot - Execute/Confirm/Elicit)

a	s
execute(hyp)	$score_{ASR} \geq 0.34$ AND $score_{NLU} \geq 0.56$
elicit	$score_{ASR} < 0.06$
confirm	<i>otherwise</i>

4.1 Experimental Results

For our experiments, we varied the amount of noise such that we match the ASR WER of the MovieBot corpus. This was achieved when SNR was set to 2.8. Figure 2 shows the simulation results (see Appendix C for learned parameters). The Y-axis in the four subfigures is the average reward per turn, the average discounted return, the average number of turns per dialog, and the execution success rate of each technique, respectively. The X-axis in all four figures shows the number of learning steps. Shaded areas highlight the standard error of the mean. RL approaches have the same dialog episodes with the user with fewer clarifying questions, and as such result in around 40% relative gain in average rewards per turn over the best rule-based method (Figure 2a). Furthermore, they achieve the same success rate (Figure 2d) with 2 fewer turns per dialog on average (Figure 2c).

To better understand the advantage of RL approaches, Figure 3 illustrates what a user experiences in order to obtain information in terms of number of turns needed to execute for each agent, where 1 means execute without any clarifying question. This figure illustrates that RL approaches learn to ask fewer clarifying questions while achieving a similar success rate. In other words, the percentage of unsuccessful executions stays the same with 6.2% for the Fixed Policy, 7.2% ($\pm 1.3\%$) for DQN, and 5.0% ($\pm 1.7\%$) for Dueling DDQN, with 0.2 less turns per execution.

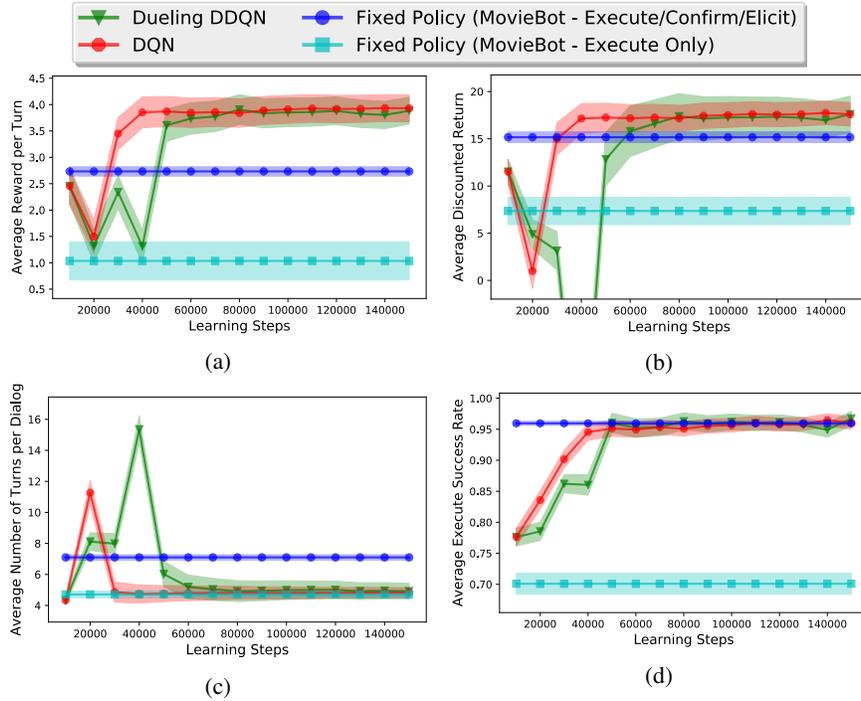


Figure 2: Results during policy learning process. a) Average reward per turn, b) average discounted return, c) average number of turns per dialog, and d) execution success rate. Shaded areas highlight the standard error of the mean. RL approaches achieve our goal of recovering from ASR/NLU errors without frustrating the user and result in around 40% relative gain in average rewards per turn over the best rule-based method, with the same success rate and 2 fewer turns per dialog.

To gain insight into the learned policy’s properties, we conducted a qualitative analysis to compare 100 dialogs generated by Dueling DDQN and Fixed Policy agents. To achieve a fair comparison, we fixed the random number generator so that the user simulator generated identical sequences of intents and utterances to interact with both agents. Overall, 28% of the dialogs were identical. Considering the remaining dialogs, two main reasons for the differences can be identified. First, the DDQN agent can decide not to clarify based on the combination of specific intent and confidence scores. Table 6 shows one dialog turn between the simulator and the two agents, along with the ground truth of the simulated utterance and intent (first column of the table), the hypothesis from MovieBot’s ASR/NLU component (second column), and the confidence scores of the hypotheses with a range between 0 and 1 (third column). In this example, since the ASR score is relatively low (0.239), the

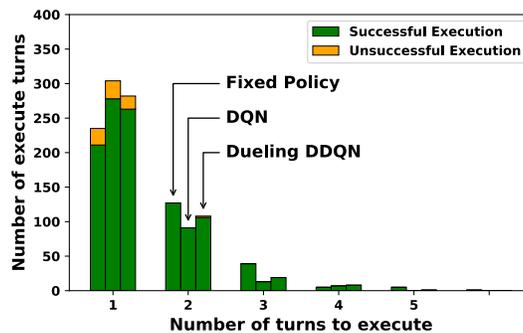


Figure 3: The number of turns needed in order to execute for Fixed Policy (Execute/Confirm/Elicit), DQN and Dueling DDQN when testing on the learned policies. RL approaches learn to ask fewer clarifying questions while achieving the same success rate.

Table 6: One dialog turn between a simulated user and Fixed Policy (MovieBot - Execute/Confirm/Elicit) agent, and Dueling DDQN agent, along with the ground truth of the utterance and intent (first column of the table), the hypothesis from the ASR/NLU components (second column), and the confidence scores of the hypotheses with a range between 0 to 1 (third column). The DDQN agent can decide not to clarify based on the combination of specific intent and confidence scores.

Fixed Policy	USER: recommend a popular movie BOT: Do you want popular movies? <confirm>
Dueling DDQN	USER: recommend a popular movie BOT: A popular movie is It, a Drama Horror film, starring Bill Skarsgård. <execute>
	reference hypothesis confidence score
	recommend a popular movie popular movies 0.239
	GetPopularMoviesIntent GetPopularMoviesIntent 0.769

Fixed Policy agent decides to ask a clarifying question (denoted as <confirm>). However, the low ASR score comes from the misrecognition of words which have small effect on the recognition of the intent as indicated by the high NLU confidence score. As such, the DDQN agent learns that the ASR confidence score is less important in this case and the agent executes the command and provides the requested information. Secondly, the DDQN agent learns that elicitation brings more information than confirmation since the user simulator repeats the original utterance. 82% of the clarifying questions asked by the DDQN agent are elicitation, as opposed to 61% for the Fixed Policy.

5 Conclusion

In this paper, we demonstrate a method for learning robust dialog policies despite upstream NLU/ASR errors caused by noisy environments. The user simulator presented here is able to mimic realistic conversations, nearly indistinguishable from human dialogs, in the context of a real-world application – the MovieBot Alexa Skill. While other simulators inject noise via stochastic processes at the intent or slot level [18], our simulator is able to leverage audio signals and white noise to train policies in a production-like environment. Using this simulator, we apply deep RL to learn dialog policies which achieve the same success rate with fewer dialog turns compared to fixed rule-based policies in noisy environments. Based on feedback from the participants in our user study, an interesting area of future work is to use sentiment to further improve the user simulator. That is, user sentiment can be sampled to generate polarized target utterances to simulate user frustration or satisfaction. Furthermore, we plan to investigate other dimensions of dialog to improve our simulator and RL-based dialog policy learning methods. These include per-interaction user satisfaction and efficiency of interactions to model the reward function. Lastly, we plan to explore the impact of using different kinds of noise on the learned policies.

Acknowledgments

We thank Jason Papis and other members of the Conversational AI group at Amazon for helpful discussions and comments. We also thank our anonymous reviewers for insightful feedback.

Appendix

A Intent Generation Hyper-parameters

	RNN Model	RNN Model + Bot Response	GRU Model	GRU Model + Bot Response
Number of Hidden Nodes per Layer	30	90	10	10
Intent and Slot Embedding Size	50	50	70	70
Number of Epochs	100	30	10	10

B Example of simulated dialogs identified as human dialogs

Bigram Simulator Dialog		RNN Simulator Dialog	
User	ummm comedy	User	ummm comedy
Bot	I'd recommend Kingsman: The Golden Circle, a comedy action film starring Taron Egerton	Bot	I think you might like Cars 3, a comedy adventure film, starring Owen Wilson.
User	Look for something else.	User	How long is the movie?
Bot	I really liked Moana, a comedy adventure film starring Auli'i Cravalho.	Bot	The movie duration is 1 hour and 42 minutes.
User	What's my favorite genre?	User	Who wrote the soundtrack?
Bot	Sci-fi, because it has been in 25 percent of your genre-related requests.	Bot	I don't know how to handle that yet. You can also say, can children watch this movie.
User	Give me a blooper from the movie inception.	User	What is a popular romantic comedy movie?
Bot	Here is an interesting inception goof. *gives goof*	Bot	I'd recommend Blind, a romance drama film starring Demi Moore.
User	How does John Wick rate?	User	Tell me about it.
Bot	John Wick has a rating of 7.3 on IMDb., based on about 300000 votes.	Bot	Here is a summary of blind from IMDb. *gives plot summary*
User	Cancel	User	That is enough.
Bot	Thank you for using MovieBot.	Bot	Thank you for using MovieBot.

C Learning parameters for DQN and Dueling DDQN for SNR of 2.8

	DQN	Dueling DDQN
Number of Hidden Layers	2	3
Number of Hidden Nodes per Layer	32	128
Embedding Size	5	30
Dropout	0.5	0.0
Learning Rate	0.0001	0.00001
Experience Replay Size	10,000	15,000
Window Size	2	8
Discount Factor	0.97	0.97
Target Model Update Interval	12,000	8,000

References

- [1] Amazon Polly. <https://aws.amazon.com/polly>
- [2] Bergstra, J.; Yamins, D.; Cox, D. D. (2013). Hyperopt: A Python Library for Optimizing the Hyperparameters of Machine Learning Algorithms, *Computational Science & Discovery*, 8(1). <https://github.com/hyperopt/hyperopt>.
- [3] Chandramohan, S.; Geist, M.; Lefevre, F.; Pietquin, O. (2011). "User Simulation in Dialogue Systems using Inverse Reinforcement Learning," in *Proceedings of Interspeech*.
- [4] Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio. (2014). On the properties of neural machine translation: Encoder-decoder approaches. In *Proceedings of the Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*.
- [5] Crook, P. A.; Marin, A. (2017). Sequence to Sequence Modeling for User Simulation in Dialog Systems. In *Proceedings of the 18th Annual Conference of the International Speech Communication Association (INTERSPEECH 2017)*, 1706–1710.
- [6] Cuayahuitl, H.; Renals, S.; Lemon, O.; Shimodaira, H. (2005). Human-computer dialogue simulation using hidden Markov models. In *Automatic Speech Recognition and Understanding, 2005 IEEE Workshop on* (pp. 290-295).
- [7] Cuayahuitl, H. (2016). SimpleDS: A Simple Deep Reinforcement Learning Dialogue System, in *Proceedings of International Workshop on Spoken Dialogue Systems (IWSDS)*.

- [8] Eckert, W.; Levin, E.; Pieraccini, R. (1997). User modeling for spoken dialogue system evaluation. In *Automatic Speech Recognition and Understanding, 1997. Proceedings., 1997 IEEE Workshop on* (pp. 80-87). IEEE.
- [9] El Asri, L.; He, J.; Suleman, K. (2016). A Sequence-to-Sequence Model for User Simulation in Spoken Dialogue Systems, *Interspeech*, 1151–1155.
- [10] Fatemi, M.; Asri, L. E.; Schulz, H.; He, J.; Suleman, K. (2016). Policy networks with two-stage training for dialogue systems, in *the Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*.
- [11] Frampton, M.; Lemon, O. (2006). Learning more effective dialogue strategies using limited dialogue move features. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, 185–192.
- [12] Georgila, K.; Henderson, J.; Lemon, O. (2006). User Simulation for Spoken Dialogue Systems: Learning and Evaluation, in *Proceedings of Interspeech*.
- [13] Georgila, K.; Traum, D. R. (2011). Reinforcement learning of argumentation dialogue policies in negotiation, in *INTERSPEECH*, 2073–2076.
- [14] Kumar, A.; Gupta, A.; Chan, J.; Tucker, S.; Hoffmeister, B.; Dreyer, M. (2017). Just ASK: Building an Architecture for Extensible Self-Service Spoken Language Understanding. arXiv:1711.00549 [cs.CL]
- [15] Lee, S.; Eskenazi, M. (2012). Pomdp-based let’s go system for spoken dialog challenge, in *Proceedings of Spoken Language Technology Workshop (SLT)*, IEEE, 61–66.
- [16] Levin, E.; Pieraccini, R.; Eckert, W. (2000). A Stochastic Model of Human-Machine Interaction for Learning Dialog Strategies, *IEEE Transactions on Speech and Audio Processing*, 1, 11–23.
- [17] Li, J.; Monroe, W.; Ritter, A.; Galley, M.; Gao, J.; Jurafsky, D. (2016). Deep Reinforcement Learning for Dialogue Generation, in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [18] Li, X.; Lipton, Z. C.; Dhingra, B.; Li, L.; Gao, J.; Chen, Y. N. (2016). A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.
- [19] Li, X.; Chen, Y. N.; Li, L.; Gao, J.; Celikyilmaz, A. (2017). End-to-End Task-Completion Neural Dialogue Systems, in *Proceedings of The 8th International Joint Conference on Natural Language Processing*.
- [20] Mahrt, T. (2016). PyAcoustics. <https://github.com/timmahrt/pyAcoustics>.
- [21] Mikolov, T.; Karafiat, M.; Burget, L.; Cernocky, J.; Khudanpur, S. (2010). Recurrent neural network based language model, in *Proceedings of INTERSPEECH*, vol. 2010, no. 9. International Speech Communication Association, 1045–1048.
- [22] Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.; Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality, in *Proceedings of NIPS*.
- [23] Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; Riedmiller, M. (2013). Playing atari with deep reinforcement learning, in *Proceedings of NIPS Deep Learning Workshop*.
- [24] Mnih, V.; Kavukcuoglu, D.; Silver, A. A.; Rusu, J.; Veness, M. G.; Bellemare, A.; Graves, M.; Riedmiller, A. K.; Fidjeland, G.; Ostro-vski, S.; Petersen, C.; Beattie, A.; Sadik, I.; Antonoglou, H.; King, D.; Kumaran, D.; Wierstra, D.; Legg, S.; Hassabis, D. (2015). Human-level control through deep reinforcement learning, *Nature*, 518(7540), 529–533.
- [25] MovieBot. <https://www.amazon.com/Amazon-MovieBot/dp/B01MRKGF5W>
- [26] Singh, S.; Litman, D.; Kearns, M.; Walker, M. (2002). Optimizing dialogue management with reinforcement learning: Experiments with the njfun system, *Journal of Artificial Intelligence Research*, 105–133.
- [27] Schatzmann, J.; Weilhammer, K.; Stuttle, M.; Young, S. (2006). A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *The knowledge engineering review*, 21(2), 97–126.
- [28] Schatzmann, J.; Young, S. (2009). The hidden agenda user simulation model. *IEEE transactions on audio, speech, and language processing*, 17(4), 733–747.
- [29] Scheffler, K.; Young, S. (2002). Automatic learning of dialogue strategy using dialogue simulation and reinforcement learning. In *Proceedings of the second international conference on Human Language Technology Research*, Morgan Kaufmann Publishers Inc, 12–18.
- [30] Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; Pineau, J. (2016). Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models, in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 3776–3783.
- [31] Sutskever, I.; Vinyals, O.; Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112.

- [32] van Hasselt, H.; Guez, A.; Silver, D. (2015). Deep Reinforcement Learning with Double Q-Learning, in *Proceedings of AAAI 2016*, arXiv:1509.06461.
- [33] Wang, Z.; Schaul, T; Hessel, M.; van Hasselt, H.; Lanctot, M.; de Freitas, N. (2016). Dueling Network Architectures for Deep Reinforcement Learning, in *Proceedings of the 33rd International Conference on Machine Learning*, vol 48.
- [34] Williams, J. D.; Young, S. (2007). Partially observable markov decision processes for spoken dialog systems, *Computer Speech & Language*, 21(2), 393–422.
- [35] Young, S.; Gasic, M.; Thomson, B.; and Williams, J. D. (2013). POMDP-based statistical spoken dialog systems: A review, *IEEE*, 101(5), 1160–1179.
- [36] Zhao, T.; Eskenazi, M. (2016). Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning, in *Proceedings of the Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*.