

# Improving Interaction Quality Estimation with BiLSTMs and the Impact on Dialogue Policy Learning

Stefan Ultes

Daimler AG

Sindelfingen, Germany

stefan.ultes@daimler.com

## Abstract

Learning suitable and well-performing dialogue behaviour in statistical spoken dialogue systems has been in the focus of research for many years. While most work which is based on reinforcement learning employs an objective measure like task success for modelling the reward signal, we use a reward based on user satisfaction estimation. We propose a novel estimator and show that it outperforms all previous estimators while learning temporal dependencies implicitly. Furthermore, we apply this novel user satisfaction estimation model live in simulated experiments where the satisfaction estimation model is trained on one domain and applied in many other domains which cover a similar task. We show that applying this model results in higher estimated satisfaction, similar task success rates and a higher robustness to noise.

## 1 Introduction

One prominent way of modelling the decision-making component of a spoken dialogue system (SDS) is to use (partially observable) Markov decision processes ((PO)MDPs) (Lemon and Pietquin, 2012; Young et al., 2013). There, reinforcement learning (RL) (Sutton and Barto, 1998) is applied to find the optimal system behaviour represented by the policy  $\pi$ . Task-oriented dialogue systems model the reward  $r$ , used to guide the learning process, traditionally with task success as the principal reward component (Gašić and Young, 2014; Lemon and Pietquin, 2007; Daubigney et al., 2012; Levin and Pieraccini, 1997; Young et al., 2013; Su et al., 2015, 2016).

An alternative approach proposes user satisfaction as the main reward component (Ultes et al., 2017a). However, the applied statistical user satisfaction estimator heavily relies on handcrafted temporal features. Furthermore, the impact of the

estimation performance on the resulting dialogue policy remains unclear.

In this work, we propose a novel LSTM-based user satisfaction reward estimator that is able to learn the temporal dependencies implicitly and compare the performance of the resulting dialogue policy with the initially used estimator.

Optimising the dialogue behaviour to increase user satisfaction instead of task success has multiple advantages:

1. The user satisfaction is more domain-independent as it can be linked to interaction phenomena independent of the underlying task (Ultes et al., 2017a).
2. User satisfaction is favourable over task success as it represents more accurately the user’s view and thus whether the user is likely to use the system again in the future. Task success has only been used as it has shown to correlate well with user satisfaction (Williams and Young, 2004).

Based on previous work by Ultes et al. (2017a), the interaction quality (IQ)—a less subjective version of user satisfaction<sup>1</sup>—will be used for estimating the reward. The estimation model is thus based on domain-independent, interaction-related features which do not have any information available about the goal of the dialogue. This allows the reward estimator to be applicable for learning in unseen domains.

The originally applied IQ estimator heavily relies on handcrafted temporal features. In this work, we will present a deep learning-based IQ estimator that utilises the capabilities of recurrent neural networks to get rid of all handcrafted fea-

<sup>1</sup>The relation of US and IQ has been closely investigated in (Schmitt and Ultes, 2015; Ultes et al., 2013).

tures that encode temporal effects. By that, these temporal dependencies may be learned instead.

The applied RL framework is shown in Figure 1. Within this setup, both IQ estimators are used for learning dialogue policies in several domains to analyse their impact on general dialogue performance metrics.

The remainder of the paper is organised as follows: in Section 2, related work is presented focusing on dialogue learning and the type of reward that is applied. In Section 3, the interaction quality is presented and how it is used in the reward model. The deep learning-based interaction quality estimator proposed in this work is then described in detail in Section 4 followed by the experiments and results both of the estimator itself and the resulting dialogue policies in Section 5.

## 2 Relevant Related Work

Most of previous work on dialogue policy learning focuses on employing task success as the main reward signal (Gašić and Young, 2014; Gašić et al., 2014; Lemon and Pietquin, 2007; Daubigney et al., 2012; Levin and Pieraccini, 1997; Young et al., 2013; Su et al., 2015, 2016). However, task success is usually only computable for pre-defined tasks e.g., through interactions with simulated or recruited users, where the underlying goal is known in advance. To overcome this, the required information can be requested directly from users at the end of each dialogue (Gašić et al., 2013). However, this can be intrusive, and users may not always cooperate.

An alternative is to use a task success estimator (El Asri et al., 2014b; Su et al., 2015, 2016). With the right choice of features, these can also be applied to new and unseen domains (Vandyke et al., 2015). However, these models still attempt to estimate completion of the underlying task, whereas our model evaluates the overall user experience.

In this paper, we show that an interaction quality reward estimator trained on dialogues from a bus information system will result in well-performing dialogues both in terms of success rate and user satisfaction on five other domains, while only using interaction-related, domain-independent information, i.e., not knowing anything about the task of the domain.

Others have previously introduced user satisfaction into the reward (Walker et al., 1998;

Walker, 2000; Rieser and Lemon, 2008b,a) by using the PARADISE framework (Walker et al., 1997). However, PARADISE relies on the existence of explicit task success information which is usually hard to obtain.

Furthermore, to derive user ratings within that framework, users have to answer a questionnaire which is usually not feasible in real world settings. To overcome this, PARADISE has been used in conjunction with expert judges instead (El Asri et al., 2012, 2013) to enable unintrusive acquisition of dialogues. However, the problem of mapping the results of the questionnaire to a scalar reward value still exists.

Therefore, we use interaction quality (Section 3) in this work because it uses scalar values applied by experts and only uses task-independent features that are easy to derive.

## 3 Interaction Quality Reward Estimation

In this work, the reward estimator is based on the interaction quality (IQ) (Schmitt and Ultes, 2015) for learning information-seeking dialogue policies. IQ represents a less subjective variant of user satisfaction: instead of being acquired from users directly, experts annotate pre-recorded dialogues to avoid the large variance that is often encountered when users rate their dialogues directly (Schmitt and Ultes, 2015).

IQ is defined on a five-point scale from five (satisfied) down to one (extremely unsatisfied). To derive a reward from this value, the equation

$$R_{IQ} = T \cdot (-1) + (iq - 1) \cdot 5 \quad (1)$$

is used where  $R_{IQ}$  describes the final reward. It is applied to the final turn of the dialogue of length  $T$  with a final IQ value of  $iq$ . A per-turn penalty of  $-1$  is added to the dialogue outcome. This results in a reward range of 19 down to  $-T$  which is consistent with related work (Gašić and Young, 2014; Vandyke et al., 2015; Su et al., 2016, e.g.) in which binary task success (TS) was used to define the reward as:

$$R_{TS} = T \cdot (-1) + \mathbb{1}_{TS} \cdot 20, \quad (2)$$

where  $\mathbb{1}_{TS} = 1$  only if the dialogue was successful,  $\mathbb{1}_{TS} = 0$  otherwise.  $R_{TS}$  will be used as a baseline.

The problem of estimating IQ has been cast as a classification problem where the target classes

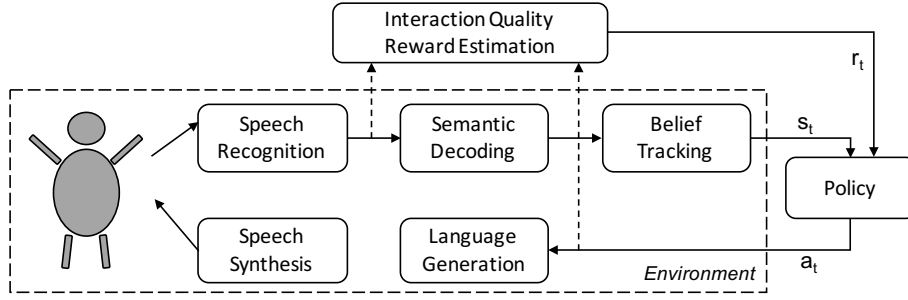


Figure 1: The RL framework integrating an interaction quality reward estimator as proposed by [Ultes et al. \(2017a\)](#). The policy learns to take action  $a_t$  at time  $t$  while being in state  $s_t$  and receiving reward  $r_t$ .




Figure 2: Modelling of temporal information in the interaction parameters used as input to the IQ estimator.

are the distinct IQ values. The input consists of domain-independent variables called interaction parameters. These parameters incorporate information from the automatic speech recognition (ASR) output and the preceding system action. Most previous approaches used this information, which is available at every turn, to compute temporal features by taking sums, means or counts from the turn-based information for a window of the last 3 system-user-exchanges<sup>2</sup> and the complete dialogue (see Fig. 2). The baseline IQ estimation approach as applied by [Ultes et al. \(2017a\)](#) (and originating from [Ultes et al. \(2015\)](#)) used a feature set of 16 parameters as shown in Table 1 with a support vector machine (SVM) ([Vapnik, 1995](#); [Chang and Lin, 2011](#)).

The LEGO corpus ([Schmitt et al., 2012](#)) provides data for training and testing and consists of 200 dialogues (4,885 turns) from the Let’s Go bus information system ([Raux et al., 2006](#)). There, users with real needs were able to call the system to get information about the bus schedule. Each turn of these 200 dialogues has been annotated with IQ (representing the quality of the dialogue up to the current turn) by three experts. The final IQ label has been assigned using the median of the three individual labels.

Previous work has used the LEGO corpus with

<sup>2</sup>a system turn followed by a user turn

Table 1: The parameters used for IQ estimation extracted on the exchange level from each user input plus counts, sums and rates for the whole dialogue (#,% ,Mean) and for a window of the last 3 turns ({.}).

	Parameter	Description
Exchange level	ASRRognitionStatus	ASR status: <i>success, no match, no input</i>
	ASRConfidence	confidence of top ASR results
	RePrompt?	is the system question the same as in the previous turn?
	ActivityType	general type of system action: <i>statement, question</i>
	Confirmation?	is system action confirm?
Dialogue level	MeanASRConfidence	mean ASR confidence if ASR is success
	#Exchanges	number of exchanges (turns)
	#ASRSuccess	count of ASR status is success
	%ASRSuccess	rate of ASR status is success
	#ASRRejections	count of ASR status is reject
%ASRRejections	rate of ASR status is reject	
Window level	{Mean}ASRConfidence	mean ASR confidence if ASR is success
	{#}ASRSuccess	count of ASR is success
	{#}ASRRejections	count of ASR status is reject
	{#}RePrompts	count of times RePrompt? is true
	{#}SystemQuestions	count of ActivityType is question

a full IQ feature set (which includes additional partly domain-related information) achieving an unweighted average recall<sup>3</sup> (UAR) of 0.55 using ordinal regression ([El Asri et al., 2014a](#)), 0.53 using a two-level SVM approach ([Ultes and Minker, 2013](#)), and 0.51 using a hybrid-HMM ([Ultes and Minker, 2014](#)). Human performance on the same task is 0.69 UAR ([Schmitt and Ultes, 2015](#)). A deep learning approach using only non-temporal features achieved an UAR of 0.55 ([Rach et al., 2017](#)).

<sup>3</sup>UAR is the arithmetic average of all class-wise recalls.

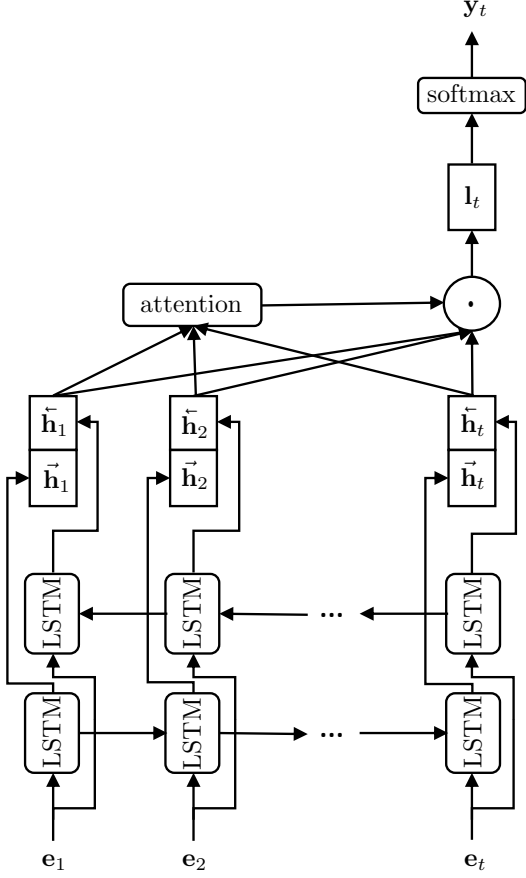


Figure 3: The architecture of the proposed BiLSTM model with self attention. For each time  $t$ , the exchange level parameter of all exchanges  $e_i$  of the sub-dialogue  $i \in \{1 \dots t\}$  are encoded to their respective hidden representation  $\mathbf{h}_i$  and are considered and weighted with the self attention mechanism to finally estimate the IQ value  $\mathbf{y}_t$  at time  $t$ .

#### 4 LSTM-based Interaction Quality Estimation

The proposed IQ estimation model will be used as a reward estimator as depicted in Figure 1. With parameters that are collected from the dialogue system modules for each time step  $t$ , the reward estimator derives the reward  $r_t$  that is used for learning the dialogue policy  $\pi$ .

The architecture of our proposed IQ estimation model is shown in Figure 3. It is based on the idea that the temporal information that has previously been explicitly encoded with the window and dialogue interaction parameter levels may be learned instead by using recurrent neural networks. Thus, only the exchange level parameters  $e_t$  are considered (see Table 1). Long Short-Term Memory (LSTM) cells are at the core of the model and have originally been proposed by Hochreiter

and Schmidhuber (1997) as a recurrent variant that remedies the vanishing gradient problem (Bengio et al., 1994).

As shown in Figure 3, the exchange level parameters form the input vector  $e_t$  for each time step or turn  $t$  to a bi-directional LSTM (Graves et al., 2013) layer. The input vector  $e_t$  encodes the nominal parameters ASRRecognitionStatus, Activity-Type, and Confirmation? as 1-hot representations. In the BiLSTM layer, two hidden states are computed:  $\vec{\mathbf{h}}_t$  constitutes the forward pass through the current sub-dialogue and  $\overleftarrow{\mathbf{h}}_t$  the backwards pass:

$$\vec{\mathbf{h}}_t = \text{LSTM}(e_t, \vec{\mathbf{h}}_{t-1}) \quad (3)$$

$$\overleftarrow{\mathbf{h}}_t = \text{LSTM}(e_t, \overleftarrow{\mathbf{h}}_{t+1}) \quad (4)$$

The final hidden layer is then computed by concatenating both hidden states:

$$\mathbf{h}_t = [\vec{\mathbf{h}}_t, \overleftarrow{\mathbf{h}}_t] \quad (5)$$

Even though information from all time steps may contribute to the final IQ value, not all time steps may be equally important. Thus, an attention mechanism (Vaswani et al., 2017) is used that evaluates the importance of each time step  $t'$  for estimating the IQ value at time  $t$  by calculating a weight vector  $\alpha_{t,t'}$ .

$$\mathbf{g}_{t,t'} = \tanh(\mathbf{h}_t^T \mathbf{W}_t + \mathbf{h}_{t'}^T \mathbf{W}_{t'} + \mathbf{b}_t) \quad (6)$$

$$\alpha_{t,t'} = \text{softmax}(\sigma(\mathbf{W}_a \mathbf{g}_{t,t'} + \mathbf{b}_a)) \quad (7)$$

$$\mathbf{l}_t = \sum_{t'} \alpha_{t,t'} \mathbf{h}_{t'} \quad (8)$$

Zheng et al. (2018) describe this as follows: “The attention-focused hidden state representation  $\mathbf{l}_t$  of an [exchange] at time step  $t$  is given by the weighted summation of the hidden state representation  $\mathbf{h}_{t'}$  of all [exchanges] at time steps  $t'$ , and their similarity  $\alpha_{t,t'}$  to the hidden state representation  $\mathbf{h}_t$  of the current [exchange]. Essentially,  $\mathbf{l}_t$  dictates how much to attend to an [exchange] at any time step conditioned on their neighbourhood context.”

To calculate the final estimate  $\mathbf{y}_t$  of the current IQ value at time  $t$ , a softmax layer is introduced:

$$\mathbf{y}_t = \text{softmax}(\mathbf{l}_t) \quad (9)$$

For estimating the interaction quality using a BiLSTM, the proposed architecture frames the task as a classification problem where each sequence is labelled with one IQ value. Thus, for

Table 2: Performance of the proposed LSTM-based variants with the traditional cross-validation setup. Due to overlapping sub-dialogues in the train and test sets, the performance of the LSTM-based models achieve unrealistically high performance.

	UAR	$\kappa$	$\rho$	eA	Ep.
LSTM	0.78	0.85	0.91	<b>0.99</b>	101
BiLSTM	<b>0.78</b>	<b>0.85</b>	<b>0.92</b>	<b>0.99</b>	100
LSTM+att	0.74	0.82	0.91	<b>0.99</b>	101
BiLSTM+att	0.75	0.83	0.91	<b>0.99</b>	93
Rach et al. (2017)	0.55	0.68	0.83	0.94	-
Ultes et al. (2015)	0.55	-	-	0.89	-

each time step  $t$ , the IQ value needs to be estimated for the corresponding sub-dialogue consisting of all exchanges from the beginning up to  $t$ . Framing the problem like this is necessary to allow for the application of a BiLSTM-approach and still be able to only use information that would be present at the current time step  $t$  in an ongoing dialogue interaction.

To analyse the influence of the BiLSTM, a model with a single forward-LSTM layer is also investigated where

$$\mathbf{h}_t = \vec{\mathbf{h}}_t. \quad (10)$$

Similarly, a model without attention is also analysed where

$$\mathbf{l}_t = \mathbf{h}_t. \quad (11)$$

## 5 Experiments and Results

The proposed BiLSTM IQ estimator is both trained and evaluated on the LEGO corpus and applied within the IQ reward estimation framework (Fig. 1) on several domains within a simulated environment.

### 5.1 Interaction Quality Estimation

To evaluate the proposed BiLSTM model with attention (BiLSTM+att), it is compared with three of its own variants: a BiLSTM without attention (BiLSTM) as well as a single forward-LSTM layer with attention (LSTM+att) and without attention (LSTM). Additional baselines are defined by Rach et al. (2017) who already proposed an LSTM-based architecture that only uses non-temporal features, and the SVM-based estimation model as originally used for reward estimation by Ultes et al. (2015).

The deep neural net models have been implemented with Keras (Chollet et al., 2015) using

Table 3: Performance of the proposed LSTM-based variants with the dialogue-wise cross-validation setup. The models by Rach et al. (2017) and Ultes et al. (2015) have been re-implemented. The BiLSTM with attention mechanism performs best in all evaluation metrics.

	UAR	$\kappa$	$\rho$	eA	Ep.
LSTM	0.51	0.63	0.78	0.93	8
BiLSTM	0.53	0.63	0.78	0.93	8
LSTM+att	0.52	0.63	0.79	0.92	40
BiLSTM+att	<b>0.54</b>	<b>0.65</b>	<b>0.81</b>	<b>0.94</b>	40
Rach et al. (2017)	0.45	0.58	0.79	0.88	82
Ultes et al. (2015)	0.44	0.53	0.69	0.86	-

the self-attention implementation as provided by Zheng et al. (2018)<sup>4</sup>. All models were trained against cross-entropy loss using RmsProp (Tieleman and Hinton, 2012) optimisation with a learning rate of 0.001 and a mini-batch size of 16.

As evaluation measures, the unweighted average recall (UAR)—the arithmetic average of all class-wise recalls—, a linearly weighted version of Cohen’s  $\kappa$ , and Spearman’s  $\rho$  are used. As missing the correct estimated IQ value by only one has little impact for modelling the reward, a measure we call the extended accuracy (eA) is used where neighbouring values are taken into account as well.

All experiments were conducted with the LEGO corpus (Schmitt et al., 2012) in a 10-fold cross-validation setup for a total of 100 epochs per fold. The results are presented in Table 2. Due to the way the task is framed (one label for each sub-dialogue), memorising effects may be observed with the traditional cross-validation setup that has been used in previous work. Hence, the results in Table 2 show very high performance, which is likely to further increase with ongoing training. However, the corresponding models are likely to generalise poorly.

To alleviate this, a dialogue-wise cross-validation setup has been employed also consisting of 10 folds of disjoint sets of dialogues. By that, it can be guaranteed that there are no overlapping sub-dialogues in the training and test sets. All results of these experiments are presented in Table 3 with the absolute improvement of the two main measures UAR and eA over the SVM-based approach of Ultes et al. (2015) visualised in Figure 4.

<sup>4</sup>Code freely available at <https://github.com/CyberZHG/keras-self-attention>




Figure 4: Absolute improvement of the IQ estimation models over the originally employed model by (Ultes et al., 2017a) for IQ-based reward estimation with the dialogue-wise cross-validation setup. UAR and eA take values from 0 to 1

The proposed BiLSTM+att model outperforms existing models and the baselines in all four performance measures by achieving an UAR of 0.54 and an eA of 0.94 after 40 epochs. Furthermore, both the BiLSTM and the attention mechanism by themselves improve the performance in terms of UAR. Based on this findings, the BiLSTM+att model is selected as reward estimator for the experiments in the dialogue policy learning setup as shown in Figure 1.

## 5.2 Dialogue Policy Learning

To analyse the impact of the IQ reward estimator on the resulting dialogue policy, experiments are conducted comparing three different reward models. The two baselines are in accordance to Ultes et al. (2017a): having the objective task success as principal reward component ( $R_{TS}$ ) and having the interaction quality estimated by a support vector machine as principal reward component ( $R_{IQ}^s$ ). TS can be computed by comparing the outcome of each dialogue with the pre-defined goal. Of course, this is only possible in simulation and when evaluating with paid subjects. This goal information is not available to the IQ estimators, nor is it required. Both baselines are compared to our proposed BiLSTM model to estimate the interaction quality used as principal reward component ( $R_{IQ}^{bi}$ ).

For learning the dialogue behaviour, a policy model based on the GP-SARSA algorithm (Gašić and Young, 2014) is used. This is a value-based method that uses a Gaussian process to approximate the state-value function. As it takes into account the uncertainty of the approximation, it

Table 4: Statistics of the domains the IQ reward estimator is trained on (LetsGo) and applied to (rest).

Domain	Code	# constraints	# DB items
LetsGo		4	-
CamRestaurants	CR	3	110
CamHotels	CH	5	33
SFRestaurants	SR	6	271
SFHotels	SH	6	182
Laptops	L	6	126

is very sample efficient and may even be used to learn a policy directly through real human interaction (Gašić et al., 2013).

The decisions of the policy are based on a summary space representation of the dialogue state tracker. In this work, the focus tracker (Henderson et al., 2014)—an effective rule-based tracker—is used. For each dialogue decision, the policy chooses exactly one summary action out of a set of summary actions which are based on general dialogue acts like *request*, *confirm* or *inform*. The exact number of system actions varies for the domains and ranges from 16 to 25.

To measure the dialogue performance, the task success rate (TSR) and the average interaction quality (AIQ) are measured: the TSR represents the ratio of dialogues for which the system was able to provide the correct result. AIQ is calculated based on the estimated IQ values of the respective model ( $AIQ^{bi}$  for the BiLSTM and  $AIQ^s$  for the SVM) at the end of each dialogue. As there are two IQ estimators, a distinction is made between  $AIQ^s$  and  $AIQ^{bi}$ . Additionally, the average dialogue length (ADL) is reported.

For the simulation experiments, the performance of the trained policies on five different domains was evaluated: Cambridge Hotels and Restaurants, San Francisco Hotels and Restaurants, and Laptops. The complexity of each domain is shown in Table 4 and compared to the LetsGo domain (the domain the estimators have been trained on).

The dialogues were created using the publicly available spoken dialogue system toolkit PyDial (Ultes et al., 2017b)<sup>5</sup> which contains an implementation of the agenda-based user simulator (Schatzmann and Young, 2009) with an additional error model. The error model simulates the required semantic error rate (SER) caused in the real system by the noisy speech channel. For each

<sup>5</sup>Code freely available at <http://www.pydial.org>




Figure 5: Results of the simulated experiments for all domains showing task success rate (TSR) only. Each value is computed after 100 evaluation / 1,000 training dialogues averaged over three trials. Numerical results with significance indicators are shown in Table 5.

domain, all three reward models are compared on three SERs: 0%, 15%, and 30%. More specifically, the applied evaluation environments are based on Env. 1, Env. 3, and Env. 6, respectively, as defined by Casanueva et al. (2017). Hence, for each domain and for each SER, policies have been trained using 1,000 dialogues followed by an evaluation step of 100 dialogues. The task success rates in Figure 5 with exact numbers shown in Table 5 were computed based on the evaluation step averaged over three train/evaluation cycles with different random seeds.

As already known from the experiments conducted by Ultes et al. (2017a), the results of the SVM IQ reward estimator show similar results in terms of TSR for  $R_{IQ}^s$  and  $R_{TS}$  in all domains for an SER of 0%. This finding is even stronger when comparing  $R_{IQ}^{bi}$  and  $R_{TS}$ . These high TSRs are achieved while having the dialogues of both IQ-based models result in higher AIQ values compared to  $R_{TS}$  throughout the experiments. Of course, only the IQ-based model is aware of the IQ concept and indeed is trained to optimise it.

For higher SERs, the TSRs lightly degrade for the IQ-based reward estimators. However, there seems to be a tendency that the TSR for  $R_{IQ}^{bi}$  is more robust against noise compared to  $R_{IQ}^s$  while still resulting in better AIQ values.

Finally, even though the differences are mostly not significant, there is also a tendency for  $R_{IQ}^{bi}$  to result in shorter dialogues compared to both  $R_{IQ}^s$  and  $R_{TS}$ .

## 6 Discussion

One of the major questions of this work addresses the impact of an IQ reward estimator on the resulting dialogues where the IQ estimator achieves better performance than previous ones. Analysing the results of the dialogue policy learning experiment leads to the conclusion that the policy learned with  $R_{IQ}^{bi}$  performs similar or better than  $R_{IQ}^s$  throughout all experiments while still achieving better average user satisfaction compared to  $R_{TS}$ . Especially for noisy environments, the improvement is relevant.

The BiLSTM clearly performs better on the LEGO corpus while learning the temporal dependencies instead of using handcrafted ones. However, it entails the risk that these learned temporal dependencies are too specific to the original data so that the model does not generalise well anymore. This would mean that it would be less suitable to be applied to dialogue policy learning for different domains. Luckily, the experiments clearly show that this is not the case.

Obviously, the experiments have only been conducted in a simulated environment and not verified in a user study with real humans. However, the general framework of applying an IQ reward estimator for learning a dialogue policy has already been successfully validated with real user experiments by Ultes et al. (2017a) and it seems rather unlikely that the changes we induce by changing the reward estimator lead to a fundamentally different result.

Table 5: Results of the simulated experiments for all domains showing task success rate (TSR), average interaction quality estimated with the SVM ( $AIQ^s$ ) and the BiLSTM ( $AIQ^{bi}$ ), and average dialogue length (ADL) in number of turns. Each value is computed after 100 evaluation / 1,000 training dialogues averaged over three trials with different random seeds. <sup>1,2,3</sup> marks statistically significant difference compared to  $R_{TS}$ , to  $R_{IQ}^s$ , and to  $AIQ^{bi}$ , respectively ( $p < 0.05$ , T-test for TSR and ADL, Mann-Whitney-U test for AIQ).

Domain	SER	TSR			$AIQ^s$		$AIQ^{bi}$		ADL		
		$R_{TS}$	$R_{IQ}^s$	$R_{IQ}^{bi}$	$R_{TS}$	$R_{IQ}^s$	$R_{TS}$	$R_{IQ}^{bi}$	$R_{TS}$	$R_{IQ}^s$	$R_{IQ}^{bi}$
CR	0%	<b>1.00</b> <sup>2,3</sup>	0.99 <sup>1</sup>	0.99 <sup>1</sup>	3.64 <sup>2</sup>	<b>3.90</b> <sup>1</sup>	3.68 <sup>3</sup>	<b>3.83</b> <sup>1</sup>	4.68	4.88	<b>4.59</b>
	15%	<b>0.97</b>	0.94	0.96	3.35 <sup>2</sup>	<b>3.65</b> <sup>1</sup>	3.45 <sup>3</sup>	<b>3.63</b> <sup>1</sup>	5.85 <sup>3</sup>	5.33	<b>5.10</b> <sup>1</sup>
	30%	<b>0.94</b>	0.92	0.90	3.15 <sup>2</sup>	<b>3.34</b> <sup>1</sup>	3.22	<b>3.30</b>	6.34	6.30	<b>6.25</b>
CH	0%	0.98	0.99	<b>0.99</b>	3.26 <sup>2</sup>	<b>3.62</b> <sup>1</sup>	3.33	<b>3.44</b>	5.71	5.61	<b>5.40</b>
	15%	<b>0.96</b> <sup>2</sup>	0.89 <sup>1,3</sup>	0.93 <sup>2</sup>	<b>2.90</b>	2.88	<b>3.14</b>	3.14	<b>6.28</b> <sup>2</sup>	7.26 <sup>1,3</sup>	6.31 <sup>2</sup>
	30%	0.86	<b>0.88</b>	0.87	2.38 <sup>2</sup>	<b>2.79</b> <sup>1</sup>	2.79 <sup>3</sup>	<b>3.02</b> <sup>1</sup>	7.94 <sup>3</sup>	7.31	<b>6.99</b> <sup>1</sup>
SR	0%	<b>0.98</b>	0.97	0.98	3.04 <sup>2</sup>	<b>3.53</b> <sup>1</sup>	3.13 <sup>3</sup>	<b>3.37</b> <sup>1</sup>	6.26	6.03	<b>5.80</b>
	15%	<b>0.90</b> <sup>3</sup>	0.88	0.84 <sup>1</sup>	2.40 <sup>2</sup>	<b>3.00</b> <sup>1</sup>	2.85 <sup>3</sup>	<b>3.01</b> <sup>1</sup>	7.99	7.55	<b>7.33</b>
	30%	0.71	0.77	<b>0.78</b>	2.03 <sup>2</sup>	<b>2.52</b> <sup>1</sup>	2.46 <sup>3</sup>	<b>2.78</b> <sup>1</sup>	9.77 <sup>3</sup>	9.41	<b>8.50</b> <sup>1</sup>
SH	0%	0.97	<b>0.99</b>	0.98	3.15 <sup>2</sup>	<b>3.52</b> <sup>1</sup>	3.17 <sup>3</sup>	<b>3.36</b> <sup>1</sup>	5.99 <sup>2</sup>	<b>5.50</b> <sup>1</sup>	5.76
	15%	0.88	0.88	<b>0.89</b>	2.63 <sup>2</sup>	<b>2.94</b> <sup>1</sup>	2.77 <sup>3</sup>	<b>3.17</b> <sup>1</sup>	7.98 <sup>3</sup>	7.59 <sup>3</sup>	<b>6.63</b> <sup>1,2</sup>
	30%	<b>0.83</b> <sup>2</sup>	0.76 <sup>1</sup>	0.80	2.50	<b>2.63</b>	2.70 <sup>3</sup>	<b>2.87</b> <sup>1</sup>	8.38	9.21	<b>8.37</b>
L	0%	0.98	<b>0.99</b>	<b>0.99</b>	3.26 <sup>2</sup>	<b>3.61</b> <sup>1</sup>	3.28	<b>3.41</b>	5.78	<b>5.44</b>	5.60
	15%	0.89	0.88	<b>0.92</b>	2.58 <sup>2</sup>	<b>2.97</b> <sup>1</sup>	2.92 <sup>3</sup>	<b>3.17</b> <sup>1</sup>	7.19	7.34	<b>6.73</b>
	30%	<b>0.80</b>	0.74	0.77	2.43	<b>2.57</b>	2.79	<b>2.92</b>	8.22 <sup>2</sup>	9.32 <sup>1,3</sup>	<b>7.97</b> <sup>2</sup>
All	0%	0.98	0.98	<b>0.98</b>	3.23 <sup>2</sup>	<b>3.65</b> <sup>1</sup>	3.31	<b>3.48</b>	5.76	5.50	<b>5.47</b>
	15%	<b>0.92</b>	0.89	0.91	2.76 <sup>2</sup>	<b>3.10</b> <sup>1</sup>	3.02 <sup>2,0</sup>	<b>3.20</b> <sup>1</sup>	7.13	7.06	<b>6.52</b>
	30%	<b>0.83</b>	0.81	0.82	2.49	<b>2.80</b>	2.78	<b>2.97</b>	8.20 <sup>2</sup>	8.23 <sup>1,3</sup>	<b>7.66</b> <sup>2</sup>

## 7 Conclusion

In this work we proposed a novel model for interaction quality estimation based on BiLSTMs with attention mechanism that clearly outperformed the baseline while learning all temporal dependencies implicitly. Furthermore, we analysed the impact of the performance increase on learned policies that use this interaction quality estimator as the principal reward component. The dialogues of the proposed interaction quality estimator show a slightly higher robustness towards noise and shorter dialogues while still yielding good performance in terms of both of task success rate and (estimated) user satisfaction. This has been demonstrated by training the reward estimator on a bus information domain and applying it to learn dialogue policies in five different domains (Cambridge restaurants and hotels, San Francisco restaurants and hotels, Laptops) in a simulated experiment.

For future work, we aim at extending the interaction quality estimator by incorporating domain-independent linguistic data to further improve the estimation performance. Furthermore, the effects of using a user satisfaction-based reward estimator needs to be applied to more complex tasks.

## References

- Yoshua Bengio, Patrice Simard, Paolo Frasconi, et al. 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166.
- Iñigo Casanueva, Paweł Budzianowski, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Stefan Ultes, Lina Rojas-Barahona, Steve Young, and Milica Gašić. 2017. A benchmarking environment for reinforcement learning based task oriented dialogue management. In *Deep Reinforcement Learning Symposium, 31st Conference on Neural Information Processing Systems (NIPS)*.
- Chih-Chung Chang and Chih-Jen Lin. 2011. LIB-SVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- François Chollet et al. 2015. Keras. <https://keras.io>.
- Lucie Daubigny, Matthieu Geist, and Olivier Pietquin. 2012. Off-policy Learning in Large-scale POMDP-based Dialogue Systems. In *Proceedings of the 37th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2012)*, pages 4989–4992, Kyoto (Japan). IEEE.
- Layla El Asri, Hatim Khouzaimi, Romain Laroche, and Olivier Pietquin. 2014a. Ordinal regression for in-



- teraction quality prediction. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3245–3249. IEEE.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2012. Reward Function Learning for Dialogue Management. In *Proceedings of the 6th Starting AI Researchers’ Symposium (STAIRS)*, pages 95–106. IOS Press.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2013. Reward shaping for statistical optimisation of dialogue management. In *Statistical Language and Speech Processing*, pages 93–101. Springer.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2014b. Task completion transfer learning for reward inference. *Proc of MLIS*.
- Milica Gašić, Catherine Breslin, Matthew Henderson, Dongho Kim, Martin Szummer, Blaise Thomson, Pirros Tsiakoulis, and Steve J. Young. 2013. On-line policy optimisation of Bayesian spoken dialogue systems via human interaction. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8367–8371. IEEE.
- Milica Gašić, Dongho Kim, Pirros Tsiakoulis, Catherine Breslin, Matthew Henderson, Martin Szummer, Blaise Thomson, and Steve J. Young. 2014. Incremental on-line adaptation of POMDP-based dialogue managers to extended domains. In *Proceedings of the 15th International Conference on Spoken Language Processing (INTERSPEECH)*, pages 140–144. ISCA.
- Milica Gašić and Steve J. Young. 2014. Gaussian processes for POMDP-based dialogue manager optimization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(1):28–40.
- Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. 2013. Hybrid speech recognition with deep bidirectional lstm. In *2013 IEEE workshop on automatic speech recognition and understanding*, pages 273–278. IEEE.
- Matthew Henderson, Blaise Thomson, and Jason Williams. 2014. The second dialog state tracking challenge. In *15th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, volume 263.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Oliver Lemon and Olivier Pietquin. 2007. Machine learning for spoken dialogue systems. In *European Conference on Speech Communication and Technologies (Interspeech’07)*, pages 2685–2688.
- Oliver Lemon and Olivier Pietquin. 2012. *Data-Driven Methods for Adaptive Spoken Dialogue Systems*. Springer New York.
- Esther Levin and Roberto Pieraccini. 1997. A stochastic model of computer-human interaction for learning dialogue strategies. In *Eurospeech*, volume 97, pages 1883–1886.
- Niklas Rach, Wolfgang Minker, and Stefan Ultes. 2017. Interaction quality estimation using long short-term memories. In *Proceedings of the 18th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 164–169. Association for Computational Linguistics.
- Antoine Raux, Dan Bohus, Brian Langner, Alan W. Black, and Maxine Eskenazi. 2006. Doing research on a deployed spoken dialogue system: One year of let’s go! experience. In *Proc. of the International Conference on Speech and Language Processing (ICSLP)*.
- Verena Rieser and Oliver Lemon. 2008a. Automatic learning and evaluation of user-centered objective functions for dialogue system optimisation. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC’08)*, pages 2356–2361, Marrakech, Morocco. European Language Resources Association (ELRA). [Http://www.lrec-conf.org/proceedings/lrec2008/](http://www.lrec-conf.org/proceedings/lrec2008/).
- Verena Rieser and Oliver Lemon. 2008b. Learning effective multimodal dialogue strategies from wizard-of-oz data: Bootstrapping and evaluation. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 638–646. ACL.
- Jost Schatzmann and Steve J. Young. 2009. The hidden agenda user simulation model. *Audio, Speech, and Language Processing, IEEE Transactions on*, 17(4):733–747.
- Alexander Schmitt and Stefan Ultes. 2015. [Interaction quality: Assessing the quality of ongoing spoken dialog interaction by experts—and how it relates to user satisfaction](#). *Speech Communication*, 74:12–36.
- Alexander Schmitt, Stefan Ultes, and Wolfgang Minker. 2012. A parameterized and annotated spoken dialog corpus of the cmu let’s go bus information system. In *International Conference on Language Resources and Evaluation (LREC)*, pages 3369–337.
- Pei-Hao Su, M. Gašić, N. Mrkšić, L. Rojas-Barahona, Stefan Ultes, D. Vandyke, T. H. Wen, and S. Young. 2016. On-line active reward learning for policy optimisation in spoken dialogue systems. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2431–2441. Association for Computational Linguistics.
- Pei-Hao Su, David Vandyke, Milica Gašić, Dongho Kim, Nikola Mrkšić, Tsung-Hsien Wen, and Steve J. Young. 2015. Learning from real users: Rating dialogue success with neural networks for reinforcement learning in spoken dialogue systems. In *Interspeech*, pages 2007–2011. ISCA.

- Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*, 1st edition. MIT Press, Cambridge, MA, USA.
- T. Tieleman and G. Hinton. 2012. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSEERA: Neural Networks for Machine Learning.
- Stefan Ultes, Paweł Budzianowski, Iñigo Casanueva, Nikola Mrkšić, Lina Rojas-Barahona, Pei-Hao Su, Tsung-Hsien Wen, Milica Gašić, and Steve Young. 2017a. Domain-independent user satisfaction reward estimation for dialogue policy learning. In *Interspeech*, pages 1721–1725. ISCA.
- Stefan Ultes, Matthias Kraus, Alexander Schmitt, and Wolfgang Minker. 2015. Quality-adaptive spoken dialogue initiative selection and implications on reward modelling. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 374–383. ACL.
- Stefan Ultes and Wolfgang Minker. 2013. Improving interaction quality recognition using error correction. In *Proceedings of the 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 122–126. Association for Computational Linguistics.
- Stefan Ultes and Wolfgang Minker. 2014. Interaction Quality Estimation in Spoken Dialogue Systems Using Hybrid-HMMs. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 208–217. Association for Computational Linguistics.
- Stefan Ultes, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gašić, and Steve J. Young. 2017b. Pydial: A multi-domain statistical dialogue system toolkit. In *ACL Demo*. Association of Computational Linguistics.
- Stefan Ultes, Alexander Schmitt, and Wolfgang Minker. 2013. On quality ratings for spoken dialogue systems – experts vs. users. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 569–578. Association for Computational Linguistics.
- David Vandyke, Pei-Hao Su, Milica Gašić, Nikola Mrkšić, Tsung-Hsien Wen, and Steve Young. 2015. Multi-domain dialogue success classifiers for policy training. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 763–770. IEEE.
- Vladimir N. Vapnik. 1995. *The nature of statistical learning theory*. Springer-Verlag New York, Inc., New York, NY, USA.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Marilyn Walker. 2000. An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research*, 12:387–416.
- Marilyn Walker, Jeanne C Fromer, and Shrikanth S. Narayanan. 1998. Learning optimal dialogue strategies: A case study of a spoken dialogue agent for email. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 2*, pages 1345–1351. Association for Computational Linguistics.
- Marilyn Walker, Diane J. Litman, Candace A. Kamm, and Alicia Abella. 1997. PARADISE: a framework for evaluating spoken dialogue agents. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics (EACL)*, pages 271–280, Morristown, NJ, USA. Association for Computational Linguistics.
- Jason D. Williams and Steve J. Young. 2004. Characterizing task-oriented dialog using a simulated asr channel. In *Proceedings of the 8th International Conference on Spoken Language Processing (Interspeech 2004)*, pages 185–188.
- Steve J. 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.
- Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. Opentag: Open attribute value extraction from product profiles. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1049–1058. ACM.