

Shades of BLEU, Flavours of Success: The Case of MultiWOZ

Tomáš Nekvinda and Ondřej Dušek

Charles University, Faculty of Mathematics and Physics

Institute of Formal and Applied Linguistics

Prague, Czech Republic

{nekvinda, odusek}@ufal.mff.cuni.cz

Abstract

The MultiWOZ dataset (Budzianowski et al., 2018) is frequently used for benchmarking context-to-response abilities of task-oriented dialogue systems. In this work, we identify inconsistencies in data preprocessing and reporting of three corpus-based metrics used on this dataset, i.e., BLEU score and Inform & Success rates. We point out a few problems of the MultiWOZ benchmark such as unsatisfactory preprocessing, insufficient or under-specified evaluation metrics, or rigid database. We re-evaluate 7 end-to-end and 6 policy optimization models in as-fair-as-possible setups, and we show that their reported scores cannot be directly compared. To facilitate comparison of future systems, we release our stand-alone standardized evaluation scripts. We also give basic recommendations for corpus-based benchmarking in future works.

1 Introduction

While human judgements are irreplaceable in dialogue systems evaluation and using full dialogue evaluation instead of evaluating isolated responses given ground-truth contexts cannot fully measure system performance (Liu et al., 2016; Takanobu et al., 2020), corpus-based evaluation metrics, such as BLEU and corpus-based entity match and success rate (Wen et al., 2017), are still very important for model development and are often used to compare models and establish state-of-the-art. We show on the MultiWOZ benchmark (Budzianowski et al., 2018), one of the most frequently used and most challenging dialogue system datasets today, that these comparisons do not hold if several basic conditions are not met, and that these conditions are not met for most of the recent works using corpus-based evaluation on this dataset. This means the assessment of progress in terms of dialogue modeling is obscured by noise coming from differences in preprocessing or metrics implementation variants.

This paper is not a critique of the MultiWOZ benchmark or of systems evaluated on it. Instead, it is a call for consistency and increased rigor in automatic evaluation. In addition to providing the analysis and identifying problems with the benchmark and current state-of-the-art reporting, we include recommendations for consistency in corpus-based score comparisons. In particular, we advocate for: (1) using standardized implementations of metrics; (2) evaluating either on detokenized surface texts, or using standardized preprocessing and postprocessing; (3) reporting the exact scripts used for evaluation; (4) release of system outputs. We also show that there is room for additional metrics of output diversity, and we add an observation on the overlap between the dialogue goals and states in training and test sections of the MultiWOZ data.

Our work can be summarized as follows:

- We identify, list, and discuss consistency issues associated with the MultiWOZ benchmark;
- We compare and re-evaluate 13 end-to-end or policy optimization systems, using a single implementation of metrics and preprocessing;
- We release the outputs of all compared systems in a unified format and provide stand-alone standardized evaluation scripts that allow for consistent comparison of future works on this dataset;¹
- In addition to standard MultiWOZ corpus-based metrics, we evaluate all systems in terms of the diversity of their outputs.

2 Related Work

Most works on evaluation methods in dialogue response generation (Deriu et al., 2021) focus on human evaluation (Walker et al., 1997), e.g., choosing the best methodology with respect to quality

¹https://github.com/Tomiinek/MultiWOZ_Evaluation

and consistency (Santhanam and Shaikh, 2019) or robustness (Dinan et al., 2019). Recent surveys in natural language generation reflect on divergence and inconsistency in human evaluation practice (Howcroft et al., 2020; Belz et al., 2020), in a similar spirit to our examination, but on a broader scale.

Despite the availability of simulator evaluation (Schatzmann et al., 2006; Young et al., 2010; Zhu et al., 2020), corpus-based metrics have been the go-to evaluation method in end-to-end neural dialogue systems since the first implementations (Wen et al., 2017; Eric and Manning, 2017) and are a de-facto standard until today (cf. Section 3.3). There are works showing problems of corpus-based metrics: limited correlation with human judgements (Novikova et al., 2017; Takanobu et al., 2020) and mixed performance depending on human reference texts used (Freitag et al., 2020) or evaluated systems (Mathur et al., 2020). Many works aim at creating more reliable metrics (Galley et al., 2015). Recent focus is on trained neural metrics (Dziri et al., 2019; Mehri and Eskenazi, 2020), but they are not yet in wide use.

Our work is probably the closest to Post (2018)’s assessment of inconsistencies in different implementations of BLEU (Papineni et al., 2002, see Section 3.2), calling for comparability and proposing a standard implementation. To our knowledge, we are the first to evaluate the use of corpus-based metrics in dialogue systems in this fashion.

3 Benchmark and Systems

3.1 The MultiWOZ dataset

The MultiWOZ 2.0 dataset (Budzianowski et al., 2018) includes about 10k task-oriented dialogues in 7 domains (restaurants, hotels, tourist attractions, trains, taxi, hospital, police) with dialogue state and system action annotation. Larger domains (restaurants, hotels, attractions, trains) have an associated database. The data was collected via human-to-human interaction on a crowdsourcing platform using the Wizard-Of-Oz approach (Wen et al., 2017). Crowd workers were instructed with goals such as booking or finding information about a restaurant or train (see Table 1). The dataset authors provided supporting code² and baselines for dialogue state tracking (DST), context-to-text (CTR), and action-to-text generation tasks.

²<https://github.com/budzianowski/multiwoz/>

Side	Utterance
User	Can you help me find a train that will arrive in birmingham new street by 16:15?
Sys.	What day would you like to travel?
User	I would like to leave Friday from Cambridge.
Sys.	TR2716 will leave cambridge at 13:01 and arrive at birmingham at 15:44. The trip will take 163 minutes and cost 75.10 pounds. Would you like to book?
User	Yes, I would like to book for 3 people.
Sys.	I booked the train. It was 225.3 and your reference is AS5SBC8P . Do you need help with anything else?
User	Yes I also need information on queens’ college.
Sys.	It’s located in west Cambridge on silver street. The postcode is cb39et. The phone number is 01223335511. The entrance fee is 2.50 pounds. Do you need anything else?
User	No, that’s great. Thank you. Good-bye.
Sys.	Thank you for using the Cambridge TownInfo centre.

Table 1: Example multi-domain conversation from the MultiWOZ dataset (conversation ID: PMUL1266).

MultiWOZ 2.1: Eric et al. (2020) released an update with re-annotated dialogue states and added explicit system action annotation.

MultiWOZ 2.2 (Zang et al., 2020) has more fixes for state annotation in 17.3% of turns, a re-defined ontology, and canonical forms for slot values (e.g. “13:00” for “1pm”) for better DST evaluation. Additionally, it introduces slot span annotations allowing easy delexicalization, which was previously based only on string matching heuristics.

3.2 Corpus-based Metrics on MultiWOZ

All standard CTR metrics on MultiWOZ – BLEU, Inform & Success rate – are calculated on *delexicalized* texts, i.e., texts where dialogue slot values, such as venue names, are replaced by placeholders (Wen et al., 2015). While using delexicalized utterances prevents errors in venue names to affect the evaluation, it prevents the use of an interactive human evaluation, model-based evaluation metrics known from open-domain dialogue research (Gao et al., 2020), or end-to-end evaluation with user simulators such as ConvLab (Zhu et al., 2020).

BLEU (Papineni et al., 2002), originally designed for machine translation (MT) evaluation, is based on comparison of n-grams in human-written references and machine-generated hypotheses. Following Wen et al. (2017), BLEU is used to measure *fluency* of output responses where the human utter-

ances are used as the reference. Using the metric for assessing fluency of the responses is not ideal, because as opposed to the intended use of BLEU, there is only a single reference available. Moreover, the set of valid responses is arguably larger for dialogue than for MT. Liu et al. (2016) show that metrics adopted from MT correlate very weakly with human judgements in dialogue responses.

Inform & Success rates: The Inform rate relates to *informable* slots, which are attributes that allow the user to constrain database searches, e.g., restaurant location or price range. The Success rate focuses on *requestable* slots, i.e., those that can be asked by the user, e.g., phone number. Both are calculated on the level of dialogues.

Su et al. (2015) consider a dialogue to be successful if the evaluated system provided all of the requested information for an entity satisfying the user’s constraints. Following this definition, Wen et al. (2017) set aside the Match rate describing whether the entity found at the *end* of each dialogue matches the user’s goal. However, MultiWOZ dialogues include multiple interleaving domains and calculating the rates only at the end is not sufficient.

Therefore, Budzianowski et al. (2018) mark a dialogue as successful if for each domain in the user’s dialogue goal: (1) the last offered entity matches (satisfies the goal constraints), and (2) the system mentioned all requestable slots required by the user. The *Inform* rate then marks the proportion of dialogues complying to (1), *Success* rate is the proportion of fully successful dialogues.

The offered entities and mentions of requestable slots are tracked over the delexicalized responses for the whole dialogue, making use of slot placeholders. If an utterance contains a slot naming an entity, e.g., restaurant name or train ID, the current dialogue state for the corresponding domain is used to query the database and an entry is sampled from the search results. At the end of a dialogue, the recorded entities and requestable slots are compared to expected values from the dialogue goal (see Appendix A for an example). The dialogue can thus be considered unsuccessful if the system does not mention a venue name or train ID at the right turn,³ does not track the user’s search constraints, or ignores the user’s requests.

³It must in practice hit the single suitable turn because responses are generated given ground-truth dialogue context.

3.3 Systems Evaluating on MultiWOZ

We discuss performance of 13 recent systems that use CTR evaluation on MultiWOZ – 7 end-to-end and 6 policy-optimization systems, which use ground-truth dialogue states during training and inference. We include models for which we got test set predictions and systems with public code for which we managed to replicate reported results.⁴

Out of the 13 compared works, 7 only report BLEU, Inform, and Success with no other evaluation; 4 use human ratings of individual outputs, and only 2 include human evaluation on full dialogues.⁵

An important representative of the end-to-end systems is DAMD (Zhang et al., 2020b). It uses a multi-action data augmentation and multiple GRU (Cho et al., 2014) decoders. Similarly, LABES (Zhang et al., 2020a) employs a few GRU-based decoders, but it represents the dialog state as a latent variable. DoTS (Jeon and Lee, 2021) also uses GRUs, but the model makes use of a BERT encoder (Devlin et al., 2019) to get a context representation. MinTL (Lin et al., 2020) applies a diff-based approach to state updates, with backbones based on the T5 and BART models (Raffel et al., 2020; Lewis et al., 2020). UBAR is based on a fine-tuned GPT-2 model (Radford et al., 2019), similarly to AuGPT (Kulhánek et al., 2021) which uses back-translations for response augmentation, and SOLOIST (Peng et al., 2020) which makes use of machine teaching (Shukla et al., 2020). We used author-provided outputs for SOLOIST and AuGPT, author-trained checkpoints for DoTS, LABES,⁶ and UBAR, and we trained DAMD and MinTL⁷ from scratch using publicly available code. DAMD, MinTL and SOLOIST use MultiWOZ 2.0; the remaining models trained on the 2.1 version. DAMD, LABES, MinTL, and UBAR are based on the same code base and use similar evaluation scripts.

We also compared 6 policy optimization models. SFN (Mehri et al., 2019), HDNO (Wang et al., 2021), and LAVA (Lubis et al., 2020) use reinforcement learning for training. HDSA (Chen et al., 2019) uses a BERT backbone and exploits the hierarchical structure of dialog acts. MarCo (Wang

⁴We were not successful in getting code, model weights, or original predictions for other systems, such as SimpleTOD (Hosseini-Asl et al., 2020), or ARDM (Wu et al., 2021).

⁵Note that full interaction is not possible with policy optimization models unless an external DST model is applied.

⁶We were able to generate outputs for 91.66% test utterances with this checkpoint. We note this in Tables 4, 5 and 6.

⁷We were only able to reproduce the T5-small model and use it in this comparison.

Delexical.	Utterance
Original	Cafe jello gallery has a free entrance fee. The address is cafe jello gallery, 13 magdalene street and the post code is cb30af. Can i help you with anything else?
MWZ 2.2	[address] has a [entrancefee] entrance fee. The address is [name], [address] and the post code is [postcode]. Can I help you with anything else?
HDSA	[attraction_name] has a free entrance fee. The address is [attraction_address] and the post code is [attraction_postcode]. Can i help you with anything else?
DAMD	[value_name] has a [value_price] entrance fee. The address is cafe jello gallery, [value_address] and the post code is [value_postcode]. Can i help you with anything else?
AuGPT	[address] has a free entrance fee. The address is cafe jello gallery, [address] and the post code is [postcode]. Can I help you with anything else?
UniConv	[attraction_name] has a [attraction_pricerange] entrance fee. The address is [attraction_name], 13 [attraction_address] and the post code is [attraction_postcode]. Can i help you with anything else?
LAVA	[attraction_name] has a free entrance fee. The address is [attraction_name], [value.count] [attraction_address] and the post code is [restaurant_postcode]. Can i help you with anything else?

Table 2: An example utterance from the MultiWOZ dataset with different styles of delexicalization. The first row shows the non-delexicalized source response. Other styles are paired with the systems that use or introduced them.

et al., 2020) and UniConv (Le et al., 2020) generate explicit system actions in parallel with the response. We use the public predictions for LAVA and the provided pretrained models for other models. UniConv and HDNO are trained on MultiWOZ 2.1, other systems use the 2.0 version. As opposed to end-to-end models, the version affects the evaluation because the ground-truth state is supplied to the model. The comparison of these systems is thus not completely fair, but we believe that the differences are small in comparison with the differences in evaluation scripts and setups (see Section 5.2).

4 Benchmark Caveats

While MultiWOZ and the associated metrics described in Section 3 represent the state-of-the-art in corpus-based dialogue evaluation practice, the benchmark has the following limitations that researchers need to be aware of: (1) delexicalization problems – imprecise delexicalization based on string matching and varying implementations thereof (Section 4.1), (2) lack of standardized postprocessing (i.e., lexicalization methods, Section 4.2), (3) database problems, i.e., multiple surface forms of database values and no information about booking availability (Section 4.3), (4) atypical metric implementations (Section 4.4), (5) lack of diversity evaluation (Section 4.5), (6) similarity between training and test data (Section 4.6).

4.1 Preprocessing

CTR evaluation metrics used in the benchmark work with *delexicalized* texts (see Section 3.2). However, the implementation of delexicalization

provided with the dataset is limited; it only applies to some expressions, leaving other slot values lexicalized. That is why most systems use their own delexicalization methods. The original delexicalization uses placeholders consisting of the domain name and the slot name, e.g. *taxi_phone*. Recent works following DAMD (Zhang et al., 2020b) remove domain names from the placeholders and determine the active domain from changes in the predicted dialogue state or model it directly.

We identified five different delexicalization styles among the 13 systems described in Section 3.3. Table 2 shows a sample system turn for which the outputs of all the delexicalization approaches are different. This is a problem since all works use their own preprocessed data as references for BLEU computation. We checked the test set for slot placeholders and found that 70.61% of the utterances contain a slot in at least one delexicalized variant and only 17.52% responses with slots exactly match for all the systems.⁸

Moreover, preprocessing scripts of some works remove contracted verb forms or keep suffixes such as “-s”, “-ly” when delexicalizing nouns or adverbs, e.g., “moderately” becomes “[pricerange]-ly”.

4.2 Postprocessing

The MultiWOZ code base does not implement backward lexicalization of texts. Out of 12 systems for which we have the source code available, only four offer scripts for lexicalizing slot values and thus allow further in-depth evaluation.

⁸8 utterances (including the example in Table 2) are pairwise different between all 5 delexicalizations.

4.3 Database: Surface Forms and Booking

The original MultiWOZ implementation of the database performs only subtle normalization of the database search constraints, such as replacing “&” with “and”. However, the slot values can have multiple valid surface forms; e.g., “4pm” and “16:00” or “the botanical gardens at cambridge university” and “cambridge university botanic gardens” correspond to the same database entities. Database query normalization is crucial for end-to-end systems, as opposed to the policy optimization models, which use ground-truth dialogue states with normalized values. The flexibility of the database might affect the Inform & Success rates, because they are based on information about database entries *complying* with the current dialogue state.

The original database does not contain any information about booking availability, because during the data collection, crowd workers were sometimes instructed to refuse a booking at a specific time, ask for another place, etc., and accept the booking with new constraints. This brings a problem into the evaluation, because some works use the ground-truth booking information (mined from the dialogue state and system action annotations) even during evaluation, whereas others ignore it and let their systems behave randomly.

4.4 Evaluation

BLEU: The original MultiWOZ BLEU implementation internally uses a trivial tokenization splitting on whitespace. However, current models often use subword tokenization and complex detokenization to remove any redundant whitespace (Sennrich et al., 2016; Kudo and Richardson, 2018). This new-style detokenization might produce words with leading or trailing punctuation. Some works ignore this fact completely, or use an alternative BLEU implementation, including tokenization, from NLTK (Bird and Loper, 2004).

Inform & Success rate: We found two main problems here. The first one comes from random database entry sampling – if multiple entities match the dialogue state, one of them is sampled at random from the database results. The set of entries complying with the dialogue state does not have to be a subset of the ground-truth set of entries complying with a given prescribed user goal from the test set. If the database results and the ground-truth set have an imperfect overlap, the sampling may choose an entry from the difference of the two sets,

which is counted as a failure. However, if an entry from the intersection of the two sets is chosen, it counts as a match, which may lead to overestimating the system performance. Some systems bypass this by comparing the sets and accepting a dialogue as matching if the sets are intersecting, or if the offered set is a non-empty subset of the ground-truth set. However, these differences result in large variances in the rates (see Section 5).

Another problem is related to the domain-oblivious delexicalization proposed by Zhang et al. (2020b). MultiWOZ responses contain slots from multiple domains at the same time very rarely, so it is sufficient to consider a single active domain for each turn. However, some works that adopt this new delexicalization use the ground-truth active domain during evaluation. Note that true domains have to be inferred from changes in ground-truth dialogue states and system actions.

4.5 Output Diversity Metrics

The standard MultiWOZ metrics do not cover the diversity of the outputs, which can show the formulaic or repetitive nature of a system’s responses (Holtzman et al., 2020). While diversity is typically measured for non-task-oriented dialogue (Li et al., 2016), we argue that it can serve as an indicator of the naturalness of using a system over longer periods of time even in task-oriented dialogue such as MultiWOZ (Oraby et al., 2018).

4.6 Dataset folds

MultiWOZ authors split the data into train, validation, and test folds randomly. Following Lampouras and Vlachos (2016)’s analysis of train-test overlap on other datasets, we inspected the goals of all 1000 test dialogues; 174 of them are also present in the train or validation folds. The test fold does not contain any unseen slot-value pairs, and has only 12 new domain-slot-value triplets. This means that the evaluation does not really check the generalization capabilities of the systems’ state tracking, and it theoretically allows the systems to memorize the whole database and bypass it during operation, which is a rather unrealistic assumption.

5 Experiments

In this section, we work with outputs produced by all systems described in Section 3.3. We: (1) unify their responses in terms of delexicalization styles, and then compare BLEU when different

System	BLEU score		Venue comparison	Inform & Success rate		Domain source
	Delexical.	Tokenization		Venue updates	Reduced search	
DAMD	DAMD	word	intersection	name, id	✓	state change
MinTL	DAMD	sub-word	intersection	name, id	✓	state change
UBAR	DAMD	sub-word	intersection	name, id	✓	state change
SOLOIST	HDSA	sub-word	-	-	-	slot names
AuGPT	AuGPT	sub-word, NLTK	first	end	✗	predicted
LABES	DAMD	word	intersection	name, id	✓	state change
DoTS	HDSA	word	sampling	name, id	✗	slot names
MarCo	HDSA	word, NLTK	subset	name, id	✗	slot names
HDSA	HDSA	word, NLTK	subset	name, id	✗	slot names
HDNO	HDSA	word	sampling	name, id	✗	slot names
SFN	HDSA	word	sampling	name, id	✗	slot names
UniConv	UniConv	word	sampling	name, id, ref.	✗	slot names
LAVA	LAVA	word	sampling	name, id	✗	slot names

Table 3: Setups of compared systems with respect to the used delexicalization method, tokenization, and Inform & Success implementation. The “Venue comparison” column describes the method of comparing offered and goal database entries, “Venue updates” indicates when the set of database entries complying to the current state is updated, “Reduced search” reflects the database implementation that ignores other search constraints if a venue name or train ID is present, and “Domain source” describes the source of information about the active turn domain.

delexicalizations are applied, (2) evaluate Inform & Success under identical conditions,⁹ (3) evaluate diversity and discuss similarity of the responses.

5.1 Setup

We report BLEU scores for six different delexicalized references (see Table 2). Five of them are styles used in HDSA, DAMD, AuGPT, UniConv, and LAVA. The sixth is delexicalization obtained from the MultiWOZ 2.2 span annotations. To make the BLEU-based comparison as fair as possible, we normalized the raw models’ outputs. First, we remove start-of-sequence tokens, all “-s” and “-ly” strings and all “s” or “es” attached to a slot placeholder. Subsequently, we lowercase the utterances, identify slots names and map them to a unified slot name ontology. The ontology contains only 18 slot names (the original domain-aware delexicalization uses around 40 slot names). It is possible to map all the slot names used in the 6 different delexicalization styles onto it. To make a single mapping possible, the result is not lossless and reduces the finer level of detail provided by some systems. For example, slots named *departure*, *destination*, and *taxi.destination* are all replaced with the *PLACE* placeholder. Finally, we pass the utterances through Moses tokenizer and detokenizer¹⁰ (Koehn et al., 2007). To calculate BLEU, we use the SacreBLEU package¹¹ (Post, 2018), which provides an

⁹Note that we work with original authors’ predictions, published pre-trained weights, or models trained from scratch, and thus we are not able to carry out a statistical analysis for the reported numbers.

¹⁰See <https://github.com/alvations/sacremoses>

¹¹See <https://github.com/mjpost/sacrebleu>

implementation compatible with the original and is now a de-facto standard in MT (cf. Section 2).

Inform & Success rates depend on the database. Our database uses fuzzy matching for the different surface forms (see Section 4.3) using the Fuzzy-Wuzzy package¹² with a similarity threshold of 90%. We use several rules to transform time strings, venue names, food types, and venue types to canonical forms matching the entries in the database (e.g., “ten o’clock p.m.” is replaced with “22:00”).

Our implementation of the Inform & Success rates follows the definition in Section 3.2. The list of offered database entries, i.e. those complying to the current dialogue state, is updated only if a venue name or a train ID is mentioned (cf. Table 3). Following HDSA, we accept a dialogue as matching if the set of offered entries is a non-empty subset of the set of entries matching the particular dialogue goal. Active domains of turns are taken from the original slot names if possible. If slot placeholders do not include the domain name, we either use model predictions if available, or estimate the domain from changes of state predictions in subsequent turns.

To better explain differences in the reported and our scores, we provide an *optimistic* Inform & Success following differences from the original implementation found in some systems, which can potentially overestimate results. In this setting, we: (1) use the intersection entry matching instead of subset matching, (2) ignore other search constraints if a name or ID is provided, (3) use ground-truth

¹²See <https://github.com/seatgeek/fuzzywuzzy>

Delexical.	End-to-end models							Policy optimization models					
	DAMD	MinTL	UBAR	SOLOIST	AuGPT	LABES*	DoTS	MarCo	HDSA	HDNO	SFN	UniConv	LAVA
MWZ 2.2	16.4	19.4	17.6	13.6	16.8	18.9	16.8	17.3	20.7	17.8	14.1	18.1	10.8
HDSA	15.5	18.6	16.3	15.1	15.5	17.1	15.7	19.0	22.5	19.4	15.6	17.9	11.4
DAMD	16.9	20.0	17.9	14.1	16.5	18.7	16.7	17.8	21.4	18.3	14.6	18.3	11.0
AuGPT	15.8	18.6	16.7	13.2	17.0	17.9	16.6	17.1	20.4	17.7	13.5	18.0	10.5
UniConv	15.1	18.2	15.9	13.7	15.5	16.9	15.5	17.6	20.6	18.1	14.1	18.8	10.9
LAVA	15.4	18.6	16.3	15.1	15.5	17.1	15.7	19.0	22.5	19.4	15.6	17.9	11.4
Reported	16.6	19.1	17.0	16.5	17.2	18.1	15.9	19.5	23.6	19.0	16.3	19.8	12.0

Table 4: Comparison of BLEU scores. The first column denotes the delexicalization style used for creating references. The highest score is highlighted for each system separately. The last row shows BLEU scores reported by authors. “*” denotes that scores for this system are computed on a subset of 91.66% test utterances.

Metric	End-to-end models							Policy optimization models					
	DAMD	MinTL	UBAR	SOLOIST	AuGPT	LABES*	DoTS	MarCo	HDSA	HDNO	SFN	UniConv	LAVA
Inform	57.9	73.7	83.4	82.3	76.6	68.5	80.4	94.5	87.9	93.3	93.4	66.7	95.9
Inform (rep.)	76.3	80.0	95.7	85.5	91.4	78.1	86.7	92.5	82.9	92.8	82.7	84.7	97.5
Inform (opt.)	73.7	79.3	88.6	86.1	78.1	75.8	84.4	96.9	91.6	97.7	96.7	67.5	97.5
Success	47.6	65.4	70.3	72.4	60.5	58.1	68.7	87.2	79.4	83.4	82.3	58.7	93.5
Success (rep.)	60.4	72.7	81.8	72.9	72.9	67.1	74.2	77.8	68.9	83.0	72.1	76.3	94.8
Success (opt.)	63.0	71.1	75.0	76.2	62.4	65.5	74.4	89.9	83.2	90.2	87.0	60.1	95.9

Table 5: Comparison of Inform & Success. “rep.” marks authors’ reported results, “opt.” denotes results for the optimistic setting (see Section 5.1). “*” for LABES marks that scores were computed on 91.66% of the test set.

active domains.¹³ Note that (2) is more permissive with respect to the system’s state tracking as the ground-truth context used during response prediction often contains ground-truth names or IDs. These are then used for the database search even if user constraints are not predicted correctly.

5.2 Results

BLEU: Table 4 summarizes BLEU evaluation using different reference texts. We notice that using a different delexicalization might substantially change the score (up to 2% BLEU absolute). Most systems perform best on the references produced by their native delexicalization used for training. We can also see that different delexicalization styles result not only in different absolute values, but also in a different relative ordering of the systems. This shows that having a single standard delexicalization (which should always be used for model evaluation and score comparison, and preferably also during model development) is very important for any fair comparison between the models. Unlike in the case of end-to-end systems, the reported scores of the policy optimization models are higher than ours.

¹³We adopt the scripts for getting ground-truth active domains from DAMD’s code base.

Inform & Success rate: Table 5 shows our and reported numbers for Inform & Success. The corpus data, i.e. ground-truth responses and dialogue states, yield Inform 93.7% and Success of 90.9%. When evaluating in the optimistic setup, these numbers grow to 97.9% and 96.6%, respectively.

Our numbers differ from the reported scores of end-to-end models to a large degree, e.g., DAMD’s reported performance is around 20% higher for both rates. However, the optimistic setting results in much lower differences. This shows that DAMD has problems with DST, which is hidden in the optimistic setup. The original UBAR numbers are very high because some ground-truth data were used during evaluation. AuGPT reports higher rates caused by a different Inform rate computation, where the set of offered venues is obtained only at the end of the dialogue. Our scores are similar to the reported ones for SOLOIST and DoTS. UniConv has the most different rates among the policy optimization models (ca. 17% for both metrics). LAVA reports higher rates similar to ours in the optimistic setting, but the difference is small and may be caused by MultiWOZ version differences. Our rates for SFN are much higher than the reported. MarCo’s and HDSA’s difference in rates can be accounted to our more flexible database.

Measure	Ref.	End-to-end models							Policy optimization models					
		DAMd	MinTL	UBAR	SOLO.	AuGPT	LAB.*	DoTS	MarCo	HDSA	HDNO	SFN	UC	LAVA
Unique tokens	1407	212	297	478	615	608	374	411	319	259	103	188	338	176
Unique trigrams	25212	1755	2525	5238	7923	5843	3228	5162	3002	2019	315	1218	2932	708
Entropy tokens	7.21	6.12	6.19	6.40	6.45	6.62	6.22	6.48	6.27	6.16	5.46	6.03	6.46	5.50
Con. ent. bigram	3.37	1.65	1.81	2.10	2.41	2.15	1.83	2.10	1.94	1.64	0.84	1.63	1.79	1.27
MSTTR-50	0.75	0.62	0.66	0.68	0.66	0.70	0.67	0.66	0.67	0.67	0.59	0.62	0.69	0.54
Avg. turn length	14.07	14.27	14.78	13.54	18.45	12.90	14.20	14.66	16.01	14.42	14.96	14.93	14.17	13.28

Table 6: Comparison of lexical diversity measures. “Ref.” shows values for delexicalized MultiWOZ 2.2 references (see Section 3). Each system has its own column. “*” denotes that scores for this system are computed on a subset of 91.66% test utterances. SOLO., LAB., UC stand for SOLOIST, LABES, and UniConv, respectively.

5.3 Evaluating Diversity

While the scores and rates differ between the evaluated systems, the generated utterances are similar and uniform (cf. Appendix B). To further understand differences between the systems, we analyzed the diversity of their responses (see Table 6).

We compare the texts on several diversity measures, following van Miltenburg et al. (2018) and Dušek et al. (2020): number of unique output tokens and trigrams, Shannon entropy and bigram conditional entropy, mean segmental type-token ratio (MSTTR-50),¹⁴ and average output length. We used the normalized texts with unified slot ontology (see Section 5.2) for the comparison. The ground-truth responses with MultiWOZ 2.2 delexicalization were used as reference. Even though the systems use different delexicalization schemes, we can draw some conclusions from the analysis. First, all the systems use rather small vocabularies. The number of used trigrams is orders of magnitude lower compared to human-produced texts. The bigram conditional entropy is also much lower for all systems. Models which employ reinforcement-learning, i.e. HDNO, SFN, and LAVA, produce the least diverse outputs. HDNO uses only 315 trigrams, which is around 1.2% of the distinct trigrams seen in reference texts. On the other hand, AuGPT, UBAR, and DoTS seem to use a broader range of expressions. Extraordinarily diverse and long are the outputs of SOLOIST. However, they are still much more closer to other models than to the human reference.

¹⁴MSTTR measures the average type-token ratio over the output text cut into segments of equal length (50 in our case). This reduces dependency on the overall text length, which is very strong in regular type-token ratio.

6 Conclusion

The MultiWOZ benchmark is unique for its size and the inclusion of a complete database, making it possible to build end-to-end task-oriented dialogue systems. Because of its naturalness and thanks to multiple fixes and revisions of state annotations, it became very popular for dialogue state tracking. However, it still has limitations for context-to-response generation, partially because of lack of standardized preprocessing and postprocessing. Since standard, easy-to-use evaluation scripts are not available, researchers are motivated to include their own modifications. This may appear unimportant, but as we showed in our analysis of 13 systems’ outputs, it results in large differences in scores and makes any comparison or tracking of progress in this area problematic.

We contribute to the solution of this problem by releasing evaluation scripts, which allow consistent evaluation of future work. We further include the evaluation of output diversity, which adds an important aspect missing from corpus-based MultiWOZ evaluation so far.

Future work should include a manual revision of MultiWOZ 2.2 span annotation to reduce training noise and to enable fair evaluation on lexicalized outputs. More important, however, is the use of human evaluation and evaluation of full dialogues in addition to corpus-based metrics (Liu et al., 2016; Takanobu et al., 2020), which is still not standard for end-to-end dialogue systems (cf. Section 3.3).

Acknowledgements

We thank the reviewers for their kind feedback. This work was supported by the Charles University grants PRIMUS/19/SCI/10, GAUK 373921, and SVV 260 575.

References

- Anya Belz, Simon Mille, and David M. Howcroft. 2020. [Disentangling the Properties of Human Evaluation Methods: A Classification System to Support Comparability, Meta-Evaluation and Reproducibility Testing](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 183–194, Dublin, Ireland.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Wenhu Chen, Jianshu Chen, Pengda Qin, Xifeng Yan, and William Yang Wang. 2019. [Semantically conditioned dialog response generation via hierarchical disentangled self-attention](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3696–3709.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar.
- Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echevoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. [Survey on Evaluation Methods for Dialogue Systems](#). *Artificial Intelligence Review*, 54:755–810.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. [Build it Break it Fix it for Dialogue Safety: Robustness from Adversarial Human Attack](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2020. [Evaluating the State-of-the-Art of End-to-End Natural Language Generation: The E2E NLG Challenge](#). *Computer Speech & Language*, 59:123–156.
- Nouha Dziri, Ehsan Kamaloo, Kory Mathewson, and Osmar Zaiane. 2019. [Evaluating Coherence in Dialogue Systems using Entailment](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3806–3812.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Kumar Goyal, Peter Ku, and Dilek Hakkani-Tür. 2020. [Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines](#). In *LREC*, pages 422–428.
- Mihail Eric and Christopher D. Manning. 2017. [A Copy-Augmented Sequence-to-Sequence Architecture Gives Good Performance on Task-Oriented Dialogue](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 468–473, Valencia, Spain.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. [BLEU might be Guilty but References are not Innocent](#). In *Proceedings of EMNLP*, pages 61–71, Online.
- Michel Galley, Chris Brockett, Alessandro Sordani, Yangfeng Ji, Michael Auli, Chris Quirk, Margaret Mitchell, Jianfeng Gao, and Bill Dolan. 2015. [deltaBLEU: A discriminative metric for generation tasks with intrinsically diverse targets](#). In *Proceedings of ACL-IJCNLP*, pages 445–450.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. [Dialogue response ranking training with large-scale human feedback data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395, Online.
- Ari Holtzman, Jan Buys, Leo Du, Maxwell Forbes, and Yejin Choi. 2020. [The Curious Case of Neural Text Degeneration](#). In *Proceedings of ICLR*, Online.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. [A simple language model for task-oriented dialogue](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 20179–20191.
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty Years of Confusion in Human Evaluation: NLG Needs Evaluation Sheets and Standardised Definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland.

- Hyunmin Jeon and Gary Geunbae Lee. 2021. [Domain state tracking for a simplified dialogue system](#).
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium.
- Jonáš Kulháněk, Vojtěch Hudeček, Tomáš Nekvinda, and Ondřej Dušek. 2021. [Augpt: Dialogue with pre-trained language models and data augmentation](#).
- Gerasimos Lampouras and Andreas Vlachos. 2016. [Imitation learning for language generation from unaligned data](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1101–1112.
- Hung Le, Doyen Sahoo, Chenghao Liu, Nancy Chen, and Steven C.H. Hoi. 2020. [UniConv: A unified conversational neural architecture for multi-domain task-oriented dialogues](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1860–1877.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A Diversity-Promoting Objective Function for Neural Conversation Models](#). In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, CA, USA.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. [MinTL: Minimalist transfer learning for task-oriented dialogue systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3391–3405, Online.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132.
- Nurul Lubis, Christian Geisshauer, Michael Heck, Hsien-chin Lin, Marco Moresi, Carel van Niekerk, and Milica Gasic. 2020. [LAVA: Latent action spaces via variational auto-encoding for dialogue policy optimization](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 465–479, Barcelona, Spain (Online).
- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2020. [Tangled up in BLEU: Reevaluating the Evaluation of Automatic Machine Translation Evaluation Metrics](#). In *ACL*, Seattle, WA, USA.
- Shikib Mehri and Maxine Eskenazi. 2020. [USR: An Unsupervised and Reference Free Evaluation Metric for Dialog Generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Shikib Mehri, Tejas Srinivasan, and Maxine Eskenazi. 2019. [Structured fusion networks for dialog](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 165–177.
- Emiel van Miltenburg, Desmond Elliott, and Piek Vossen. 2018. [Measuring the Diversity of Automatic Image Descriptions](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1730–1741.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why We Need New Evaluation Metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2243–2253, Copenhagen, Denmark.
- Shereen Oraby, Lena Reed, Shubhangi Tandon, Sharath T.S., Stephanie Lukin, and Marilyn Walker. 2018. [Controlling Personality-Based Stylistic Variation with Neural Natural Language Generators](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 180–190.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. [Few-shot natural language generation for task-oriented dialog](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 172–182, Online.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language Models are Unsupervised Multitask Learners](#). Technical report, OpenAI.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Sashank Santhanam and Samira Shaikh. 2019. [Towards Best Experiment Design for Evaluating Dialogue System Output](#). In *Proceedings of INLG*, pages 88–94, Tokyo, Japan.
- Jost Schatzmann, Karl Weilhammer, Matt Stuttle, and Steve Young. 2006. [A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies](#). *The Knowledge Engineering Review*, 21(2):97–126.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany.
- Swadheen Shukla, Lars Liden, Shahin Shayandeh, Eslam Kamal, Jinchao Li, Matt Mazzola, Thomas Park, Baolin Peng, and Jianfeng Gao. 2020. [Conversation Learner - a machine teaching tool for building dialog managers for task-oriented dialog systems](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 343–349, Online.
- Pei-Hao Su, David Vandyke, Milica Gašić, Dongho Kim, Nikola Mrkšić, Tsung-Hsien Wen, and Steve Young. 2015. [Learning from real users: rating dialogue success with neural networks for reinforcement learning in spoken dialogue systems](#). In *Proceedings of Interspeech 2015*, pages 2007–2011.
- Ryuichi Takanobu, Qi Zhu, Jinchao Li, Baolin Peng, Jianfeng Gao, and Minlie Huang. 2020. [Is Your Goal-Oriented Dialog Model Performing Really Well? Empirical Analysis of System-wise Evaluation](#). In *SIGdial*, pages 297–310, Online.
- Marilyn A. Walker, Diane J. Litman, Candace A. Kamm, and Alicia Abella. 1997. [PARADISE: A Framework for Evaluating Spoken Dialogue Agents](#). In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, pages 271–280, Madrid, Spain.
- Jianhong Wang, Yuan Zhang, Tae-Kyun Kim, and Yunjie Gu. 2021. [Modelling hierarchical structure between dialogue policy and natural language generator with option framework for task-oriented dialogue system](#). In *Proceedings of ICLR*.
- Kai Wang, Junfeng Tian, Rui Wang, Xiaojun Quan, and Jianxing Yu. 2020. [Multi-domain dialogue acts and response co-generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7125–7134, Online.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. [Semantically Conditioned LSTM-based Natural Language Generation for Spoken Dialogue Systems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721, Lisbon, Portugal.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. [A network-based end-to-end trainable task-oriented dialogue system](#). In *Proceedings of EACL*, pages 438–449, Valencia, Spain.
- Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu. 2021. [Alternating recurrent dialog model with large-scale pre-trained language models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1292–1301, Online.
- Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. 2010. [The Hidden Information State model: A practical framework for POMDP-based spoken dialogue management](#). *Computer Speech & Language*, 24(2):150–174.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. [MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117, Online.
- Yichi Zhang, Zhijian Ou, Min Hu, and Junlan Feng. 2020a. [A probabilistic end-to-end task-oriented dialog model with latent belief states towards semi-supervised learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9207–9219.
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020b. [Task-oriented dialog systems that consider multiple appropriate responses under the same context](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9604–9611.
- Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takanobu, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020. [ConvLab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 142–149, Online.

Goal database entries (ID):	19212, 19185, 19197, 19219 (cheap and Chinese search constraints)			
Goal requested information:	ADDRESS, POST			
Turn utterances	Dialogue state	Offered Entries	Active Domain	Provided Info
<i>User:</i> I'd like to find a restaurant that serves cheap Chinese food! <i>Ground-truth system:</i> Charlie Chan is located in the centre and it is cheap! Would you like me to book it for you? <i>System:</i> there are COUNT PRICE FOOD restaurant. do you have a preference of area?	food: chinese, pricerange: cheap	no update	restaurant	∅
<i>User:</i> Yes, go ahead and book it and please send me the address. <i>Ground-truth system:</i> The address is Regent Street City Centre. What day and time would you like to book? How many people? <i>System:</i> okay, i have booked NAME, which is located at ADDRESS. is there anything else i can do for you today?	food: chinese, pricerange: cheap, area: centre	19212, 19185, 19219	restaurant	ADDRESS
<i>User:</i> I do not need to book it. What is the postcode and the area please? <i>Ground-truth system:</i> The postcode is cb21db <i>System:</i> NAME is in the AREA and the postcode is POST. is there anything else i can help you with today?	food: chinese, pricerange: cheap, area: centre	19212, 19185, 19219	restaurant	POST, ADDRESS
<i>User:</i> Thanks very much! <i>Ground-truth system:</i> You're welcome! Have a great day! Good-bye. <i>System:</i> thank you for using our services.	food: chinese, pricerange: cheap, area: centre	no update	-	POST, ADDRESS

Table 7: Inform & Success calculation on the predictions of the UBAR model (conversation ID: SNG0805).

A Inform & Success Calculation Details

Table 7 walks through the process of Inform & Success calculation. Rows group conversation turns. The first column shows the last user utterance, the corresponding ground-truth system response and the delexicalized and normalized generated response. The second column shows the current dialogue state. The “Offered entities” column shows the changes of the set of matching venue or train IDs. Note that the set is updated only if the generated response contains the NAME or TRAINID placeholder. The “Active domain” column shows the currently active domain. The “Provided Info” column lists requestable slots mentioned until the given point.

This sample conversation is (1) matching, i.e. contributes positively to the Inform rate, because the set of offered entities or database entries at the

end of the dialogue is a non-empty subset of the set of goal database entries, and also (2) successful because it is matching and all requested information defined by the dialogue goal, i.e., the restaurant address and post code, were provided. Note that the ground-truth context is used to generate system utterances during the evaluation, and thus the dialogue state might contain information from past ground-truth system utterances (such as the area slot in our example, which was never mentioned in the generated system or user utterances).

B Generated Texts Comparison

Tables 8 and 9 give a comparison of delexicalized utterances generated by all examined systems for two different contexts. Slot placeholders are shown in capital letters. While there are some differences in the systems’ behavior, the outputs often contain very similar phrasing.

Context	User: Hello, I am doing research on plant life in England, and visiting cambridge as part of that, can you give me more info on the cambridge university botanic gardens.
System	Normalized utterance
Corpus	that sounds like an exciting trip! NAME is located in the AREA on ADDRESS. its admission fee is PRICE.
DAMD	the address is ADDRESS and the postcode is POST. is there anything else i can help you with?
MinTL	NAME is a TYPE in the AREA. it is located at ADDRESS. their phone number is PHONE.
UBAR	NAME is a TYPE located in the AREA at ADDRESS. it is PRICE to enter and is located on ADDRESS. is there anything else i can help you with?
SOLOIST	certainly! NAME is located at ADDRESS. their phone number is PHONE. is there anything else i can help you with?
AuGPT	NAME is a TYPE located in the AREA of town. the entrance fee is 4 pounds. the address is ADDRESS.
LABES	NAME is located in the AREA area of town
DoTS	NAME is a TYPE in the AREA. it's free to enter. do you need any other information?
MarCO	NAME is a nightclub in the AREA of town. the address is ADDRESS and the postcode is POST.
HDSA	it is a museum in the AREA of town. the entrance fee is free. would you like the address?
HDNO	NAME is a museum in the AREA. the address is ADDRESS. the phone number is PHONE. is there anything else i can help you with?
SFN	NAME is located at ADDRESS. the postcode is POST. the phone number is PHONE. can i help you with anything else?
UniConv	sure! it is a TYPE in the AREA part of town. would you like the address?
LAVA	the NAME is located at ADDRESS, postcode POST. the phone number is PHONE.

Table 8: Side-by-side comparison of system outputs for the 1st system turn of the conversation PMUL3012.

Context	User: I'm looking for a nice place to eat. System: There are many fine places to eat. What type of cuisine would you prefer? User: I don't care about the cuisine type, but I want it to be somewhere expensive and in the centre please.
System	Normalized utterance
Corpus	there are COUNT PRICE eateries in the AREA. how about one of my favorites - NAME? they serve FOOD food and it's delicious!
DAMD	how about NAME? it is an FOOD restaurant that is in the AREA of town and is in the PRICE price range .
MinTL	there are COUNT restaurant that meet your criteria. do you have a preference for type of food?
UBAR	there are COUNT PRICE restaurant in the AREA. do you have a cuisine preference?
SOLOIST	how about NAME? it is an FOOD restaurant that is in the AREA of town and is in the PRICE price range .
AuGPT	there are 33 PRICE restaurants in the AREA. What type of food would you like?
LABES	i have COUNT restaurants that meet your criteria.
DoTS	i would recommend NAME. would you like to make a reservation?
MarCO	there are COUNT restaurant that meet your criteria. i recommend NAME.
HDSA	i recommend NAME. would you like a reservation?
HDNO	NAME is a FOOD restaurant in the AREA. would you like me to book a table for you?
SFN	there are COUNT PRICE restaurant in the AREA. what type of cuisine would you like?
UniConv	there are 33 PRICE TYPE in the AREA. what type of food are you looking for?
LAVA	NAME is located at ADDRESS, postcode POST. the phone number is PHONE.

Table 9: Side-by-side comparison of system outputs for the 2nd system turn of the conversation PMUL2489.