

# ConvLab-3: A Flexible Dialogue System Toolkit Based on a Unified Data Format

Qi Zhu<sup>\*1</sup> Christian Geishauer<sup>\*2</sup> Hsien-chin Lin<sup>\*2</sup> Carel van Niekerk<sup>\*2</sup>  
Baolin Peng<sup>3</sup> Zheng Zhang<sup>1</sup> Shutong Feng<sup>2</sup> Michael Heck<sup>2</sup> Nurul Lubis<sup>2</sup>  
Dazhen Wan<sup>1</sup> Xiaochen Zhu<sup>4</sup> Jianfeng Gao<sup>3</sup> Milica Gašić<sup>†2</sup> Minlie Huang<sup>†1</sup>

<sup>1</sup>Tsinghua University, Beijing, China

<sup>2</sup>Heinrich Heine University Düsseldorf, Düsseldorf, Germany

<sup>3</sup>Microsoft Research, Redmond, USA

<sup>4</sup>University of Cambridge, Cambridge, England

<sup>1</sup>{zhu-q18,wandz19}@mails.tsinghua.edu.cn {z-zhang,aihuang}@tsinghua.edu.cn

<sup>2</sup>{geishaus,linh,niekerk,heckmi,lubis,gasic}@hhu.de

<sup>3</sup>{bapeng,jfgao}@microsoft.com <sup>4</sup>xz479@cam.ac.uk

## Abstract

Task-oriented dialogue (TOD) systems function as digital assistants, guiding users through various tasks such as booking flights or finding restaurants. Existing toolkits for building TOD systems often fall short of in delivering comprehensive arrays of data, models, and experimental environments with a user-friendly experience. We introduce ConvLab-3: a multifaceted dialogue system toolkit crafted to bridge this gap. Our unified data format simplifies the integration of diverse datasets and models, significantly reducing complexity and cost for studying generalization and transfer. Enhanced with robust reinforcement learning (RL) tools, featuring a streamlined training process, in-depth evaluation tools, and a selection of user simulators, ConvLab-3 supports the rapid development and evaluation of robust dialogue policies. Through an extensive study, we demonstrate the efficacy of transfer learning and RL and showcase that ConvLab-3 is not only a powerful tool for seasoned researchers but also an accessible platform for newcomers<sup>1</sup>.

## 1 Introduction

Task-oriented dialogue (TOD) systems converse with their users in natural language to help them fulfil a task, such as booking a flight or finding a restaurant. Unlike chit-chat dialogues, a critical aspect of these systems is that they are grounded in an ontology that contains domains, slots, and values which describe the dialogue task, i.e. user goal, as well as including domain-specific databases.

<sup>1</sup>ConvLab-3 is publicly available at <https://github.com/ConvLab/ConvLab-3> under Apache License 2.0. The demonstrative video accompanying this paper is available at <https://youtu.be/t6HVTJCeGLo>.

<sup>\*</sup>These authors contributed equally to this work.

<sup>†</sup>These authors share the senior authorship of this work.

There are two distinct capabilities that TOD systems need to exhibit. They need to *track* the state of the dialogue and based on that *decide* on the next action to take in order to steer the conversation towards fulfilling the user’s goal (Young et al., 2007). The architecture of TOD systems typically adopts a modular approach, often encompassing components like dialogue state trackers and policies, and may include language understanding or generation units, as depicted in Figure 1. The complexity of a TOD system necessitates a toolkit with advanced, easily integrable modules allowing for straightforward training, evaluation, and combination.

The vast amount of possible user behaviours and tasks that a TOD system might assist with necessitates the study of generalization and transfer towards new users and datasets. While many datasets for studying task-oriented dialogue have been proposed (Wen et al., 2016; Mrkšić et al., 2017; Byrne et al., 2019; Eric et al., 2020; Rastogi et al., 2020; Zhu et al., 2020a; Feng et al., 2022), the various dialogue, ontology and database formats hinder researchers from validating their models on unseen data. In this work we propose a unified format to bridge the gap between different TOD datasets and models and provide a unified training and evaluation framework that accelerates the study of generalization capabilities. Once a dataset is transformed into the unified format, it can be immediately used by supported models. Similarly, once a model supports the unified format, it can access all supported datasets. This feature reduces the cost of adapting  $M$  models to  $N$  datasets from  $M \times N$  to  $M + N$ .

The dialogue policy, as the decision-making component of a TOD system, is pivotal to the success or failure of a dialogue task. It is typically optimized using reinforcement learning (RL), necessitating additional components such as algorithms,

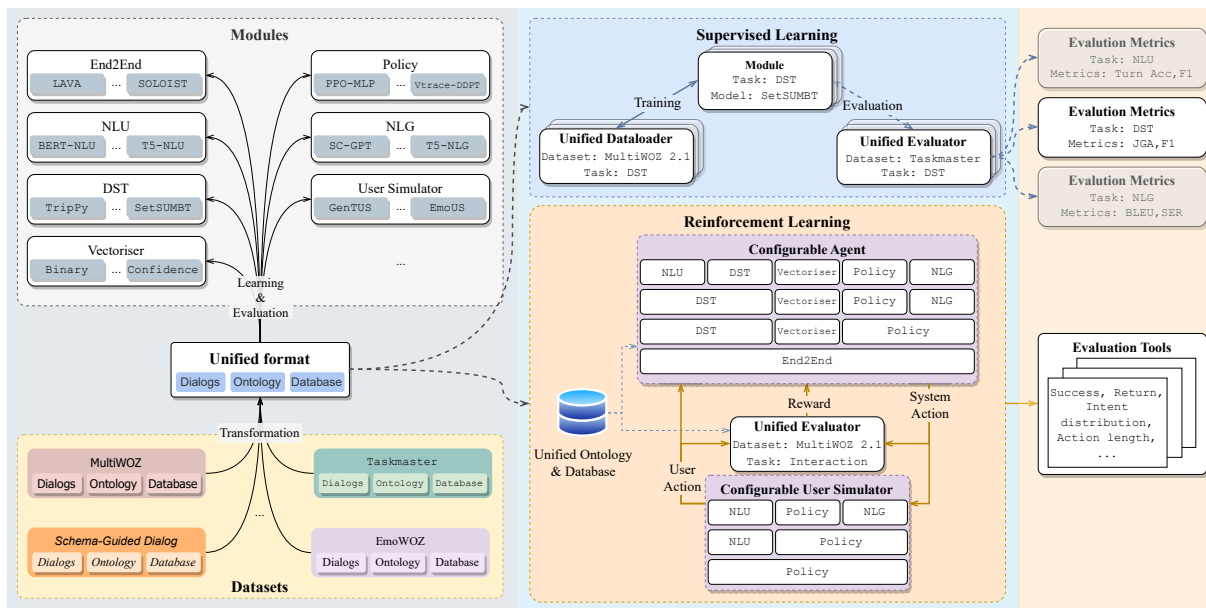


Figure 1: ConvLab-3: The unified format serves as a bridge, connecting diverse datasets and dialogue models. It streamlines the integration of various TOD modules, including supervised learning, evaluation, and a wide array of essential evaluation metrics, thanks to the unified data loader and evaluator. These modules can be incorporated, either in the agent or user simulator, through a configuration file, defining the environment for interactive evaluation and reinforcement learning.

evaluation tools, and user simulators. Realistic user simulators are essential for conducting interactive evaluations and tests against varied user behaviours, in order to accurately mirror real-world scenarios. ConvLab-3 streamlines RL-based development and assessment of dialogue policies. We achieve this by offering a configurable RL environment, evaluation tools for thorough insights, and multiple user simulators to explore generalization capabilities towards new user behaviours, as depicted in Figure 1.

ConvLab-3 is especially useful for practitioners seeking to construct a dialogue system without extensive expertise. Additionally, it provides a fast, convenient, and dependable platform for both novice and experienced researchers to conduct experiments. In particular, it enables: (1) researchers to perform experiments across a variety of datasets, (2) developers to construct an dialogue system using custom datasets, and (3) community contributors to consistently add models and datasets. In summary, our contributions are:

- A unified data format which allows for easy generalisation and transfer learning experiments across different datasets.
- A convenient RL framework and access to different user simulators, accelerating the development and evaluation of dialogue policies.

- Providing a broad collection of compatible datasets and state-of-the-art models.

## 2 Related work

While Rasa (Bocklisch et al., 2017), NeMo (Kuchaiev et al., 2019) and DialogueStudio (Zhang et al., 2023) provide unified data formats, they do not have RL tools or user simulators for interactive training and evaluation of dialogue systems. ParlAI (Miller et al., 2017) includes a *reward* attribute in their unified format, but without accessible RL tools. PyDial (Ultes et al., 2017) and the predecessors of ConvLab-3 (Lee et al., 2019; Zhu et al., 2020b) provide reinforcement learning toolkits, however they lack a unified format and thus the possibility to study generalization across datasets. Moreover, PyDial and previous versions of ConvLab do not provide multiple data-driven user simulators and their training evaluation provides no tools for in-depth analysis. In addition, none of the above toolkits provide a sufficient set of state-of-the-art models for the different components in a TOD system.

## 3 Unified Format

In our unified format, a dataset consists of (1) an **ontology** that defines the annotation schema, (2) **dialogues** with transformed annotations, and (3) a

Dataset	Dataset Annotations					
	Goal	DA-U	DA-S	State	API	Database
Camrest (2016)	✓	✓	✓	✓		✓
WOZ 2.0 (2017)		✓		✓		
KVRET (2017)		✓		✓	✓	
DailyDialog (2017)		✓				
Taskmaster-1 (2019)		✓	✓	✓		
Taskmaster-2 (2019)		✓	✓	✓		
MultiWOZ 2.1 (2020)	✓	✓	✓	✓		✓
SGD (2020)		✓	✓	✓	✓	
MetaLWOZ (2020)	✓					
CrossWOZ (2020a)	✓	✓	✓	✓	✓	✓
Taskmaster-3 (2021)		✓	✓	✓	✓	
EmoWOZ (2022)	✓	✓	✓	✓		✓

Table 1: Annotations of current unified datasets. DA-U/DA-S is dialogue acts annotation of user/system.

**database** that links to external knowledge sources (see Figure 1).

Typically converting the formats of different datasets is not straightforward, hindering format adaptation of existing and new corpora. However, in ConvLab-3 we provide detailed guidelines and scripts that make the process of format adaptation straightforward and error-free. ConvLab-3 offers a large number of datasets in the unified format as shown in Table 1, whilst also simplifying the process of adding new datasets.

Moreover, as shown in Listing 1, we provide utility functions to process the unified datasets, such as delexicalization, splitting data for few-shot learning, and loading data for specific tasks. Based on the unified format, evaluations of common tasks across models and corpora are standardized, which facilitates comparability. More details of already supported datasets and tasks can be found in Appendix A and B, respectively.

### 3.1 Ontology

Following Budzianowski et al. (2018) and Rastogi et al. (2020), an ontology consists of: (1) Domains and their slots in a hierarchical format. Each slot has a Boolean flag indicating whether it is a categorical slot (whose value set is fixed). (2) All possible intents in dialogue acts. (3) Possible dialogue acts appearing in the dialogues. Each act is comprised of intent, domain, slot, and speaker (i.e., system or user). (4) Template dialogue state. We also provide a natural language description, if available, for each domain, slot, and intent to facilitate few-shot learning (Mi et al., 2022) and domain transfer (Lin et al., 2021b).

```

from convlab.util import *

dataset_name = "multiwoz21"
# load dataset: a dict maps data_split to dialogues
dataset = load_dataset(dataset_name)
# load dataset in a predefined order with a custom
# split ratio for reproducible few-shot experiments
dataset = load_dataset(dataset_name, \
                       split2ratio={"train": 0.01})

# load ontology and database similarly
ontology = load_ontology(dataset_name)
database = load_database(dataset_name)
# query the database with domain and state
state = {"hotel": {"area": "east", \
                  "price range": "moderate"}}
res = database.query("hotel", state, topk=3)

# Example functions based on the unified format
# load the user turns in the test set for NLU task
nlu_data = load_nlu_data(dataset, "test", "user")
# dataset-agnostic delexicalization
dataset, delex_vocab = create_delex_data(dataset)

```

Listing 1: Example usage of unified datasets.

### 3.2 Dialogues

We unify the format of dialogue annotations included in many datasets and commonly used by dialogue models while keeping the original format of annotations that only appear in specific datasets. As we integrate more datasets in the future, we will expand the unified format to include more common annotations.

For a dialogue in the unified format, dialogue-level information includes the dataset name, data split (training or test), unique dialogue ID, involved domains, user goal, etc. Following MultiWOZ (Budzianowski et al., 2018), a user goal has informable slot-value pairs, requestable slots, and a natural language instruction summarizing the goal.

Turn-level information includes speaker, utterance, dialogue acts, state, database result, etc (see Appendix H for an example). Each dialogue act is a list of tuples, each tuple consisting of intent, domain, slot, and value. According to the value, we divide dialogue acts into three groups: (1) **categorical** for slots whose value set is predefined in the ontology (e.g., inform the weekday of a flight). (2) **non-categorical** for slots whose values can not be enumerated (e.g., inform the address of a hotel). (3) **binary** for intents without actual values (e.g., request the address of a hotel). The state is initialized by the template state as defined by the ontology and updated during the conversation, containing slot-value pairs of involved domains. A database result is a list of entities retrieved from the database or other knowledge sources. We list common annotations included in the unified data format and the tasks they support in Appendix B.

Other dataset-specific annotations are retained in their original formats.

### 3.3 Database/API Interface

To unify the interaction with different types of databases, we define a `BaseDatabase` class that has an abstract query function to be customized. The query function takes the current domain, dialogue state, and other custom arguments as input and returns a list of top-k candidate entities. By inheriting `BaseDatabase` and overriding the query function, we can easily access different databases/APIs and retrieve the result with a unified format.

### 3.4 Evaluation

To provide a comparable evaluation setup for all TOD tasks supported by the unified format, we provide unified evaluation scripts. These scripts include commonly used metrics such as: turn accuracy (ACC) and dialogue act F1 score for natural language understanding (NLU) (Zhu et al., 2020b), joint goal accuracy (JGA) and slot F1 score for dialogue state tracking (DST) (Li et al., 2021), BLEU and slot error rate (SER) for natural language generation (NLG) (Wen et al., 2015), BLEU and Combined score (Comb.) for End2End dialogue modeling (Mehri et al., 2019), turn accuracy, slot-value F1 score and SER for user simulators (Lin et al., 2021a, 2022).

## 4 Integrated Models

ConvLab-3 provides a wide array of standard and state-of-the-art models covering all modules in a TOD system. This allows straightforward plug-and-play experimentation when developing a specific module, as well as building TOD systems easily on custom datasets. A model is considered integrated once it implements the corresponding module interface and supports processing datasets in the unified format.

Besides existing models in ConvLab-2 (Zhu et al., 2020b), we integrate new transformer-based models supporting the unified data format, including SetSUMBT (van Niekerk et al., 2021) and TripPy (Heck et al., 2020) for dialogue state tracking (DST), DDPT (Geishauser et al., 2022) and LAVA (Lubis et al., 2020) for policy learning, SC-GPT (Peng et al., 2020) for natural language generation (NLG), and SOLOIST (Peng et al., 2021) with T5 as backbone model (Peng et al., 2022) for end-to-end modeling (End2End). We also integrate

multiple powerful data-driven user simulators (US): TUS (Lin et al., 2021a) that outputs user dialogue acts, GenTUS (Lin et al., 2022) that outputs both user dialogue acts and response, and EmoUS (Lin et al., 2023) that additionally outputs emotions.

In addition, we apply text-generation models to solve the tasks of TOD modules (see Appendix C). We provide a range of models built upon T5 (Raffel et al., 2020), covering NLU, DST, NLG, etc. We also provide an interface to instruct large language models (LLMs) such as ChatGPT and LLaMa (Touvron et al., 2023) to serve as different modules such as user simulators, NLU, DST, NLG, etc. See Heck et al. (2023) for an example of how ChatGPT can be instructed to serve as a DST model. All integrated models are shown in Appendix B.

## 5 Reinforcement Learning Toolkit

The difficulty of building a comprehensive TOD toolkit lies in the fact that it needs to support not only supervised but also reinforcement learning. As shown in Figure 1, this includes functionalities to build configurable agents and user simulators consisting of different modules, an evaluator to provide reward signals, and analysis tools to evaluate the training process and RL algorithms. ConvLab-3 supports the straightforward combination of components with an easy-to-use configuration file, including the definition of the interactive environment given by the choice of user policy and its components, see Appendix G for an example.

The dialogue policy module obtains the semantic information of the DST (and NLU) as input and produces a list of atomic actions  $[(domain_1, intent_1, slot_1), \dots]$  as output, e.g.  $[(hotel, inform, phone), (hotel, inform, addr)]$ , which results in a large action space due to the high number of possible atomic actions and their combinations. As the input is on semantic level while the policy network expects vectorized input, ConvLab-3 provides a `Vectoriser` class that acts as communication module between semantic and vector representation. We treat the `Vectoriser` as an additional pipeline module, which allows straightforward investigation of different vectorization strategies in a plug-and-play fashion. Moreover, policy networks can be used off-the-shelf while only the `Vectoriser` needs to be adapted. ConvLab-3 provides a base `Vectoriser` class that can be easily adapted, as well as common vectorization strategies. In addition, we add the possi-



bility for masking certain actions as in Ultes et al. (2017). This allows controllability of the policy output and facilitates learning during RL due to reduction of the large action space. Moreover, in addition to the on-policy RL algorithms REINFORCE (Sutton et al., 1999) and PPO (Schulman et al., 2017), which are already implemented in ConvLab-2, we provide the state-of-the-art continual RL model DDPT together with state-of-the-art algorithms VTRACE (Espeholt et al., 2018) and CLEAR (Rolnick et al., 2019) for off-policy (Sutton and Barto, 2018) and continual RL (Khetarpal et al., 2022), respectively.

## 5.1 Evaluation Tools

Understanding the policy behaviour allows researchers to fine-tune their algorithm or reward model in an informed manner to improve performance. The analysis of policy behaviour can be done by studying 1) the efficiency of actions, i.e. how many atomic actions are taken in a turn, 2) how the selected intents are distributed in a turn, 3) actual dialogue interactions. The average number of atomic actions is an important indicator of information overload, which a user simulator can handle well in contrast to humans. The intent distribution reveals policy preferences and possible exploitations of imperfect user simulators.

ConvLab-3 is the first toolkit to provide these set of measures and evaluation tools together with the common measurements of task success, return and average number of turns. Moreover, actual dialogues can be observed for in-depth evaluation.

## 6 Supervised Learning Experiments

Conducting supervised learning experiments on multiple TOD datasets is convenient with the unified data format. We believe this feature will encourage researchers to build general dialogue models that perform well on various data as well as to investigate knowledge transfer. In these experiments, we demonstrate the ease of evaluating a model’s knowledge transfer abilities using our unified format. Initially, we pre-train all models on the Schema-Guided Dialogue (SGD) (Rastogi et al., 2020) and Taskmaster-1&2&3 (Byrne et al., 2019, 2021) datasets jointly. These models are then fine-tuned on MultiWOZ 2.1 (Eric et al., 2021) in full-data or low-resource settings. To configure these different training setups, one only needs to make a few changes to the unified dataloader parameters,

DST	MultiWOZ 2.1					
	1%		10%		100%	
	JGA ↑	Slot F1 ↑	JGA ↑	Slot F1 ↑	JGA ↑	Slot F1 ↑
T5-DST	14.5 <b>22.9</b>	68.5 <b>74.9</b>	35.5 <b>41.2</b>	84.8 <b>87.1</b>	52.6 53.1	91.9 92.0
SetSUMBT	7.8 <b>22.7</b>	41.8 <b>77.2</b>	37.0 <b>43.8</b>	84.4 <b>88.2</b>	50.3 50.7	90.8 91.2
NLG	SER ↓	BLEU ↑	SER ↓	BLEU ↑	SER ↓	BLEU ↑
	19.0 <b>9.8</b>	20.2 <b>25.8</b>	6.9 <b>5.5</b>	31.3 <b>32.9</b>	3.7 3.5	35.8 35.8
	27.3 <b>9.5</b>	14.1 <b>26.3</b>	11.2 <b>6.9</b>	28.4 28.6	4.8 5.3	33.6 32.1
End2End	Comb.	BLEU	Comb.	BLEU	Comb.	BLEU
SOLOIST	19.8 <b>42.2</b>	0.4 <b>10.4</b>	48.0 <b>62.0</b>	10.0 <b>15.9</b>	67.0 <b>71.4</b>	16.8 17.5

Table 2: Comparison between models without pre-training (1st row) and with pre-training (2nd row) in both the low-resource and full-data settings.

as depicted in Listing 1. For low-resource fine-tuning, we set the data ratios of both training and validation set to 1% and 10%.

In the low-resource setting, we observe that pre-training is beneficial, as evidenced in Table 2. Specifically for the end-to-end model SOLOIST, pre-training also proves advantageous in the full-data setting. This may be attributed to the increased complexity of the end-to-end modeling task. These findings emphasize that transfer learning can be successfully implemented in ConvLab-3 in a straightforward way. This enables: (1) developers to leverage knowledge from existing datasets for application in smaller, custom settings; (2) newcomers to explore the capabilities of various models; and (3) experienced researchers to evaluate the generalisability of their proposed methods, as well as to compare them to the available state-of-the-art benchmarks. For an example of joint training across multiple datasets and retrieval based data augmentation, see Appendix D and E.

## 7 Reinforcement Learning Experiments

ConvLab-3 supports a convenient way to run RL training and evaluation supported by the unified format and availability of multiple user simulators. To showcase this, we run transfer learning experiments as well as experiments with multiple user simulators.

### 7.1 Transfer Learning

We utilize the DDPT policy model with VTRACE as algorithm and consider four different data set




Figure 2: Pre-training then RL training experiments with the DDPT model in interaction with the rule-based simulator. Shaded regions show standard error. Each model is evaluated on 9 different seeds.

US for training	US for testing		
	ABUS	TUS	GenTUS
ABUS	<b>0.93</b>	0.71	0.56
TUS	0.87	0.79	0.59
GenTUS	0.89	<b>0.86</b>	<b>0.63</b>

Table 3: The strict success rates of PPO-MLP policies trained on ABUS, TUS, and GenTUS when evaluated with various user simulators.

scenarios for supervised pre-training: (1) **scratch** that does not use pre-training, (2) **SGD** that pre-trains on SGD, (3) **1% MWOZ** that pre-trains on 1% of MultiWOZ data, and (4) **SGD->1%MWOZ** that pre-trains on SGD data and afterwards 1% of MultiWOZ data. The experiments are conducted on the semantic level, leveraging the rule-based dialogue state tracker and the rule-based user simulator (Schatzmann et al., 2007a) of ConvLab-3.

The results, depicted in Figure 2, show a similar trend for all models and metrics. Nevertheless, Figure 2(a) reveals that pre-training on SGD does not yield an advantage for the starting performance, compared to training from scratch, while it leads to better results by the end of training. Moreover, the number of actions taken in a turn and the probability of taking a request intent, shown in Figure 2(b) and (c), is initially much lower for the model trained on SGD only. This indicates that the behaviour learned from SGD differs significantly from the behaviour on MultiWOZ. Refer to Appendix F for more results and experiments. Our unique evaluation tools thus provides essential insights into both metrics and the behaviour of the agent.

## 7.2 Evaluation across Different User Simulators

To enable a policy to generalize to diverse user behaviour, it’s crucial to train and evaluate policy models across various user simulators. ConvLab-3 not only offers state-of-the-art data-driven user simulation models but also a configurable interactive environment for evaluation and reinforcement learning, as illustrated in Figure 1. In these experiments, we utilise a multi-layer perceptron (MLP) policy trained with the PPO algorithm, using three distinct user simulators: ABUS (Schatzmann et al., 2007b), TUS, and GenTUS. We then evaluate the resulting policies using each of these simulators.

The results, listed in Table 3, show that the policy trained with ABUS excels only in ABUS evaluations, while the GenTUS-trained policy outperforms others in GenTUS and TUS evaluations, but performs slightly worse than ABUS-trained policies in ABUS evaluations. This result highlights the importance of cross-US training and evaluation to show the generalizability of the dialogue policy. Conducting such experiments is made straightforward in ConvLab-3, as the user simulator model can be easily changed within the configuration file.

## 8 Conclusion

In this paper, we present the dialogue system toolkit ConvLab-3, which puts a large number of datasets under one umbrella through our proposed unified data format. The usage of the unified format facilitates comparability and significantly reduces the implementation cost required for conducting experiments on multiple datasets. In addition, we provide recent powerful models for all components of a dialogue system and provide a convenient RL toolkit which enables researchers to easily build, train, analyze and evaluate dialogue systems.

We showcase the advantages of the unified format and RL toolkit in a large number of experiments, ranging from pre-training to RL training. The release of ConvLab-3 supports the community in developing the next generation of task-oriented dialogue systems.

## 9 Limitations

As ConvLab-3 is built for text-based TOD systems, we do not currently provide support for speech. One solution for this is the usage of speech recognition and text-to-speech interfaces such as Whisper (Radford et al., 2023) and WaveNet (van den Oord et al., 2016). Secondly, while we provide several datasets in the unified format together with conversion scripts, the conversion of a new dataset still requires manual effort such as normalizing ontologies and transforming dialogue annotations. Lastly, ConvLab-3, currently, only supports the commonly used hierarchical dialogue state representation (Budzianowski et al., 2018) but not yet state representations such as the graph-based state (Andreas et al., 2020) and tree-structure state (Cheng et al., 2020). We consider these limitations as future work to further improve our toolkit.

## Acknowledgements

This work was supported by the National Key Research and Development Program of China (No. 2021ZD0113304), the National Science Foundation for Distinguished Young Scholars (with No. 62125604), the NSFC projects (Key project with No. 61936010 and regular project with No. 61876096), and sponsored by Tsinghua-Toyota Joint Research Fund.

This work was also supported by the DYMO project which has received funding from the European Research Council (ERC) provided under the Horizon 2020 research and innovation programme (Grant agreement No. STG2018 804636). In addition, it was supported by an Alexander von Humboldt Sofja Kovalevskaja Award endowed by the German Federal Ministry of Education and Research. Computational infrastructure and support were provided by the Centre for Information and Media Technology at Heinrich Heine University Düsseldorf and the Google Cloud Platform.

## References

- Jacob Andreas, John Bufe, David Burkett, Charles Chen, Josh Clausman, Jean Crawford, Kate Crim, Jordan DeLoach, Leah Dorner, Jason Eisner, Hao Fang, Alan Guo, David Hall, Kristin Hayes, Kellie Hill, Diana Ho, Wendy Iwaszuk, Smriti Jha, Dan Klein, Jayant Krishnamurthy, Theo Lanman, Percy Liang, Christopher H. Lin, Ilya Lintsbakh, Andy McGovern, Aleksandr Nisnevich, Adam Pauls, Dmitriy Petters, Brent Read, Dan Roth, Subhro Roy, Jesse Rusak, Beth Short, Div Slomin, Ben Snyder, Stephon Striplin, Yu Su, Zachary Tellman, Sam Thomson, Andrei Vorobev, Izabela Witoszko, Jason Wolfe, Abby Wray, Yuchen Zhang, and Alexander Zotov. 2020. [Task-oriented dialogue as dataflow synthesis](#). *Transactions of the Association for Computational Linguistics*, 8:556–571.
- Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. 2017. [Rasa: Open source language understanding and dialogue management](#). *arXiv preprint arXiv:1712.05181*.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Bill Byrne, Karthik Krishnamoorthi, Saravanan Ganesh, and Mihir Kale. 2021. [TicketTalk: Toward human-level performance with end-to-end, transaction-based dialog systems](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 671–680, Online. Association for Computational Linguistics.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. [Taskmaster-1: Toward a realistic and diverse dialog dataset](#). 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)*, pages 4516–4525, Hong Kong, China. Association for Computational Linguistics.
- Jianpeng Cheng, Devang Agrawal, Héctor Martínez Alonso, Shruti Bhargava, Joris Driesen, Federico Flego, Dain Kaplan, Dimitri Kartsaklis, Lin Li, Dhivya Piraviperumal, Jason D. Williams, Hong Yu, Diarmuid Ó Séaghda, and Anders Johansen. 2020. [Conversational semantic parsing for dialog state tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8107–8117, Online. Association for Computational Linguistics.

- Mihail Eric, Nicole Chartier, Behnam Hedayatnia, Karthik Gopalakrishnan, Pankaj Rajan, Yang Liu, and Dilek Hakkani-Tur. 2021. [Multi-sentence knowledge selection in open-domain dialogue](#). In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 76–86, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. [MultiWOZ 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking base-lines](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 422–428, Marseille, France. European Language Resources Association.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. [Key-value retrieval networks for task-oriented dialogue](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49, Saarbrücken, Germany. Association for Computational Linguistics.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. 2018. [IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1407–1416. PMLR.
- Shutong Feng, Nurul Lubis, Christian Geischauser, Hsien-chin Lin, Michael Heck, Carel van Niekerk, and Milica Gasic. 2022. [EmoWOZ: A large-scale corpus and labelling scheme for emotion recognition in task-oriented dialogue systems](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4096–4113, Marseille, France. European Language Resources Association.
- Christian Geischauser, Carel van Niekerk, Hsien-chin Lin, Nurul Lubis, Michael Heck, Shutong Feng, and Milica Gašić. 2022. [Dynamic dialogue policy for continual reinforcement learning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 266–284, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Michael Heck, Nurul Lubis, Benjamin Ruppik, Renato Vukovic, Shutong Feng, Christian Geischauser, Hsien-chin Lin, Carel van Niekerk, and Milica Gasic. 2023. [ChatGPT for zero-shot dialogue state tracking: A solution or an opportunity?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 936–950, Toronto, Canada. Association for Computational Linguistics.
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geischauser, Hsien-Chin Lin, Marco Moresi, and Milica Gasic. 2020. [TripPy: A triple copy strategy for value independent neural dialog state tracking](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 35–44, 1st virtual meeting. Association for Computational Linguistics.
- Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. 2022. [Towards continual reinforcement learning: A review and perspectives](#). *Journal of Artificial Intelligence Research*, 75:1401–1476.
- Oleksii Kuchaiev, Jason Li, Huyen Nguyen, Oleksii Hrinchuk, Ryan Leary, Boris Ginsburg, Samuel Kriman, Stanislav Beliaev, Vitaly Lavrukhin, Jack Cook, et al. 2019. [Nemo: a toolkit for building ai applications using neural modules](#). *arXiv preprint arXiv:1909.09577*.
- Sungjin Lee, Qi Zhu, Ryuichi Takanobu, Zheng Zhang, Yaoqin Zhang, Xiang Li, Jinchao Li, Baolin Peng, Xiujun Li, Minlie Huang, and Jianfeng Gao. 2019. [ConvLab: Multi-domain end-to-end dialog system platform](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 64–69, Florence, Italy. Association for Computational Linguistics.
- Jinchao Li, Baolin Peng, Sungjin Lee, Jianfeng Gao, Ryuichi Takanobu, Qi Zhu, Minlie Huang, Hannes Schulz, Adam Atkinson, and Mahmoud Adada. 2020. [Results of the multi-domain task-completion dialog challenge](#). In *Proceedings of the 34th AAAI Conference on Artificial Intelligence, Eighth Dialog System Technology Challenge Workshop*.
- Jinchao Li, Qi Zhu, Lingxiao Luo, Lars Liden, Kaili Huang, Shahin Shayandeh, Runze Liang, Baolin Peng, Zheng Zhang, Swadheen Shukla, Ryuichi Takanobu, Minlie Huang, and Jianfeng Gao. 2021. [Multi-domain task-oriented dialog challenge ii at dstc9](#). In *AAAI-2021 Dialog System Technology Challenge 9 Workshop*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Hsien-Chin Lin, Shutong Feng, Christian Geischauser, Nurul Lubis, Carel van Niekerk, Michael Heck, Benjamin Ruppik, Renato Vukovic, and Milica Gašić. 2023. [Emous: Simulating user emotions in task-oriented dialogues](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 2526–2531, New York, NY, USA. Association for Computing Machinery.
- Hsien-chin Lin, Christian Geischauser, Shutong Feng, Nurul Lubis, Carel van Niekerk, Michael Heck, and Milica Gasic. 2022. [GenTUS: Simulating user behaviour and language in task-oriented dialogues with](#)



- generative transformers. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 270–282, Edinburgh, UK. Association for Computational Linguistics.
- Hsien-chin Lin, Nurul Lubis, Songbo Hu, Carel van Niekerk, Christian Geischauser, Michael Heck, Shutong Feng, and Milica Gasic. 2021a. [Domain-independent user simulation with transformers for task-oriented dialog systems](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 445–456, Singapore and Online. Association for Computational Linguistics.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021b. [Leveraging slot descriptions for zero-shot cross-domain dialog StateTracking](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648, Online. Association for Computational Linguistics.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [What makes good in-context examples for GPT-3?](#) In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114, Dublin, Ireland and Online. Association for Computational Linguistics.
- Nurul Lubis, Christian Geischauser, 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). International Committee on 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, Stockholm, Sweden. Association for Computational Linguistics.
- Fei Mi, Yitong Li, Yasheng Wang, Xin Jiang, and Qun Liu. 2022. [CINS: comprehensive instruction for few-shot learning in task-oriented dialog systems](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. [ParLAI: A dialog research software platform](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 79–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. [Neural belief tracker: Data-driven dialogue state tracking](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788, Vancouver, Canada. Association for Computational Linguistics.
- Baolin Peng, Michel Galley, Pengcheng He, Chris Brockett, Lars Liden, Elnaz Nouri, Zhou Yu, Bill Dolan, and Jianfeng Gao. 2022. [Godel: Large-scale pre-training for goal-directed dialog](#). *arXiv preprint arXiv:2206.11309*.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2021. [Soloist: Building task bots at scale with transfer learning and machine teaching](#). *Transactions of the Association for Computational Linguistics*, 9:807–824.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujuan 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. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine Mcleavey, and Ilya Sutskever. 2023. [Robust speech recognition via large-scale weak supervision](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 28492–28518. PMLR.
- 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.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. 2019. [Experience replay for continual learning](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007a. [Agenda-based user simulation for bootstrapping a POMDP dialogue system](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 149–152, Rochester, New York. Association for Computational Linguistics.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007b. [Agenda-based user simulation for bootstrapping a POMDP dialogue system](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 149–152, Rochester, New York. Association for Computational Linguistics.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *arXiv preprint arXiv:1707.06347*.
- Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement learning: an introduction*, second edition. The MIT Press.
- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. [Policy gradient methods for reinforcement learning with function approximation](#). In *Advances in Neural Information Processing Systems*, volume 12. MIT Press.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- 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 Young. 2017. [PyDial: A multi-domain statistical dialogue system toolkit](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada. Association for Computational Linguistics.
- Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alexander Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. [Wavenet: A generative model for raw audio](#). In *Arxiv*.
- Carel van Niekerk, Andrey Malinin, Christian Geisshauser, Michael Heck, Hsien-chin Lin, Nurul Lubis, Shutong Feng, and Milica Gasic. 2021. [Uncertainty measures in neural belief tracking and the effects on dialogue policy performance](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7901–7914, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Wenhui Wang, Hangbo Bao, Shaohan Huang, Li Dong, and Furu Wei. 2021. [MiniLMv2: Multi-head self-attention relation distillation for compressing pre-trained transformers](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2140–2151, Online. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young. 2016. [Conditional generation and snapshot learning in neural dialogue systems](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2153–2162, Austin, Texas. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gašić, 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. Association for Computational Linguistics.
- Steve Young, Jost Schatzmann, Karl Weilhammer, and Hui Ye. 2007. [The Hidden Information State Approach to Dialog Management](#). In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing - ICASSP '07*, volume 4, pages IV–149–IV–152.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. [Glm-130b: An open bilingual pre-trained model](#). *arXiv preprint arXiv:2210.02414*.
- Jianguo Zhang, Kun Qian, Zhiwei Liu, Shelby Heinecke, Rui Meng, Ye Liu, Zhou Yu, Huan Wang, Silvio Savarese, and Caiming Xiong. 2023. [Dialogstudio: Towards richest and most diverse unified dataset collection for conversational ai](#).
- Qi Zhu, Kaili Huang, Zheng Zhang, Xiaoyan Zhu, and Minlie Huang. 2020a. [CrossWOZ: A large-scale Chinese cross-domain task-oriented dialogue dataset](#). *Transactions of the Association for Computational Linguistics*, 8:281–295.
- Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takano, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020b. [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. Association for Computational Linguistics.

## A Annotations of current unified datasets

The statistics and annotations of already supported datasets are listed in Table 4.

Dataset	Statistics			Dataset Annotations					
	#Dialogues	Avg. Turns	Domains	Goal	DA-U	DA-S	State	API Result	Database
Camrest (Wen et al., 2016)	676	10.8	1	✓	✓	✓	✓		✓
WOZ 2.0 (Mrkšić et al., 2017)	1200	7.4	1		✓		✓		
KVRET (Eric et al., 2017)	3030	5.3	3		✓		✓	✓	
DailyDialog (Li et al., 2017)	13118	7.9	10		✓				
Taskmaster-1 (Byrne et al., 2019)	13175	21.2	6		✓	✓	✓		
Taskmaster-2 (Byrne et al., 2019)	17303	16.9	7		✓	✓	✓		
MultiWOZ 2.1 (Eric et al., 2020)	10438	13.7	8	✓	✓	✓	✓		✓
Schema-Guided (Rastogi et al., 2020)	22825	20.3	45		✓	✓	✓	✓	
MetaLWOZ (Li et al., 2020)	40203	10.4	51	✓					
CrossWOZ (Zhu et al., 2020a)	6012	16.9	6	✓	✓	✓	✓	✓	✓
Taskmaster-3 (Byrne et al., 2021)	23757	20.1	1		✓	✓	✓	✓	
EmoWOZ (Feng et al., 2022)	11434	14.6	8	✓	✓	✓	✓		✓

Table 4: Statistics and annotations of current unified datasets. DA-U/DA-S is dialogue acts annotation of user/system.

## B Tasks and models supported by the unified data format

Tasks and models already supported by the unified data format are shown in Table 5.

## C Example Serialized Dialogue Acts and State

The example of serialized dialogue acts and states is shown in Table 6.

## D Joint Training

In this experiment, we investigate the effect of training a model on multiple datasets jointly instead of separately. For joint training, we merge MultiWOZ, SGD, and Taskmaster datasets into one and train a single model, which requires the model to handle datasets with different ontologies. Intuitively, the advantage of joint training is that knowledge transfer is bi-directional and persists for the whole training period, while the disadvantage is that there may be inconsistent labels for similar inputs on different datasets, potentially confusing the models.

To avoid confusion, for T5-NLU, T5-DST, and T5-NLG, we prepend the dataset name to the original input to distinguish data from different datasets. For SetSUMBT, we only predict the state of the target dataset. Since SGD may have several services for one domain, we normalize the service name to the domain name (e.g., Restaurant\_1 to Restaurant) when evaluating NLU and DST. However, similar slots of different services (e.g., city and location) will still confuse the model. While further normalization may help, we are aiming to compare independent training and joint training instead of achieving SOTA performance. For Taskmaster-1/2/3, we evaluate each sample with

the corresponding ontology and then calculate the metrics on all test samples of three datasets. In addition, on SGD and Taskmaster, we build pseudo user goals for TUS and GenTUS by accumulating constraints and requirements in user dialogue acts during conversations.

We compare independent training and joint training in Table 7. MultiWOZ, SGD, and Taskmaster have 8K, 16K, and 43K dialogues for training respectively. Joint training on these datasets does not lead to substantial performance drops in most cases, indicating that models have sufficient capacity to encode knowledge of different datasets simultaneously. However, joint training does not always improve performance either. It consistently improves the End2End model SOLOIST but makes no difference to T5-NLU. For other models, the gains vary with the dataset. Associating with the previous pre-training-then-fine-tuning experiment, we think the difference may be attributed to the varying task complexity on different datasets. When the original data of a certain dataset are sufficient for a model to solve the task, including other datasets via joint training may not bring further benefit.

## E Retrieval Augmentation

We further explore transferring knowledge from other datasets through retrieval-based data augmentation. Here we only consider the single-turn NLU task where the input is an utterance since utterance-level similarity is easier to model than dialogue-level similarity. For each utterance in the target dataset, we retrieve the top-k ( $k \in \{1, 3\}$ ) most similar utterances from other datasets measured by the MiniLMv2 model (Wang et al., 2021) using Sentence Transformers (Reimers and Gurevych, 2019). We then use retrieved samples in two ways:

Task	Input	Output	Models
RG	Context	Response	T5RG, LLMs
Goal2Dial	Goal	Dialogue	T5Goal2Dialogue
NLU	Context	DA-U	T5NLU, BERTNLU, MILU, LLMs
DST	Context	State	T5DST, (Set)SUMBT, TripPy, LLMs
Policy	State, DA-U, Database	DA-S	DDPT, PPO, PG
Word-Policy	Context, State, Database	Response	LAVA
NLG	DA-S	Response	T5NLG, SC-GPT, LLMs
End2End	Context, Database	State, Response	SOLOIST
User Simulator	Goal, DA-S	DA-U, (Response)	TUS, GenTUS, EmoUS, LLMs

Table 5: Tasks and models supported by the unified data format. RG is response generation without database support. Goal2Dial is generating a dialogue from a user goal. NLU is natural language understanding. BERTNLU, MILU, PPO, PG are from ConvLab-2 (Zhu et al., 2020b). Currently supported LLMs include LLaMA-2 (Touvron et al., 2023), ChatGLM2 (Zeng et al., 2022), and OpenAI models such as ChatGPT (through API).

---

**User:** I am looking for a *cheap* restaurant.  
**System:** Is there a particular area of town you prefer?  
**User:** In the *centre* of town.

---

**DA-U:** [inform][restaurant]([area][centre])  
**State:** [restaurant]([area][centre],[price range][cheap])  
**DA-S:** [recommend][restaurant]([name][Zizzi Cambridge])

---

**System:** I would recommend *Zizzi Cambridge*.

---

Table 6: Example serialized dialogue acts and state. Dialogue acts are in the form of “[intent] [domain] ([slot] [value],...);...”. State is in the form of “[domain] ([slot] [value],...);...”. Multiple items are separated by a semi-colon.

1. Augment training data. Models are trained on both original training data and retrieved data.
2. Additionally input the retrieved samples as in-context examples, including retrieved utterances and their dialogue acts, as shown in Table 8. Different from Liu et al. (2022), the retrieved samples are from other datasets instead of the target dataset and we will train models on augmented samples.

Since different datasets have different ontologies (i.e., definitions of intent, domain, slot), we prepend the corresponding dataset name to an input utterance as in the joint training experiment. We use the T5-NLU model and try two model sizes T5-Small and T5-Large. We fine-tune the models on MultiWOZ using the same settings as in the pre-training-then-fine-tuning experiment.

## F RL Experiments

### F.1 Transfer Learning

In this section, we provide additional plots for the experiments conducted in Section 7. Figure 3 depicts success rate (which is less strict success rate),

average number of turns as well as the average return. Moreover, we show additional intent probabilities. We can observe in Figure 3(e) that the policy pre-trained on only SGD data uses more offer intents initially, which reflects the behaviour in the data set. The probability of the offer intent then decreases whereas the probability for the recommend intent increases during learning.

### F.2 Training with Uncertainty Features

#### F.2.1 Experimental Setup

The simplified vectorizer module in ConvLab-3 makes it possible to easily include dialogue related features that might benefit dialogue policy learning. One such example is given due to the problem of resolving ambiguities in conversations. Humans naturally identify these ambiguities and resolve the uncertainty resulting from them. For a dialogue system to be robust to ambiguity, it is crucial to identify and resolve these uncertainties (van Niekerk et al., 2021).

ConvLab-3 provides the SetSUMBT dialogue belief tracker which achieves SOTA performance in terms of the accuracy of its uncertainty estimates. Using the vectorizer class to incorporate these features we can train a policy using the uncertainty features obtained from SetSUMBT. To illustrate the effectiveness of uncertainty features during RL, we train a PPO policy using these features. The template-based NLG in ConvLab-3 also allows for the inclusion of noise (van Niekerk et al., 2021) in generated responses which allows for uncertainties to arise during the conversation, simulating a more realistic conversation.

#### F.2.2 Result Analysis

Figure 4 (a) reveals that the policy trained using uncertainty features performs at least as well as the



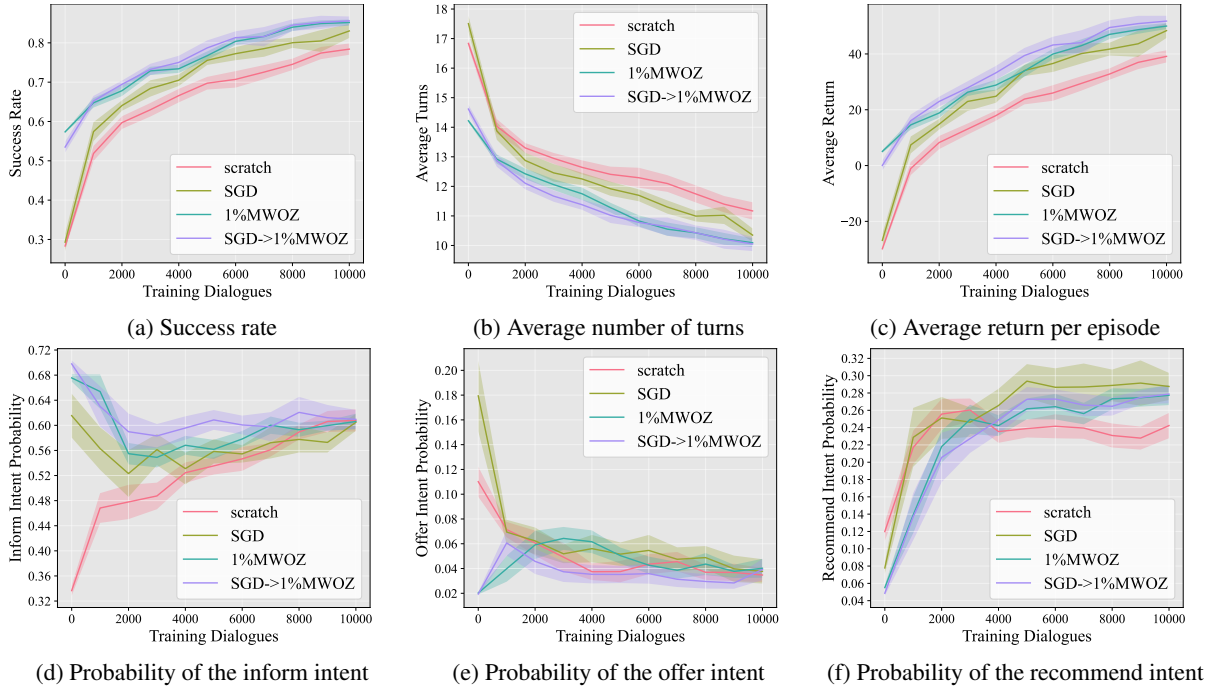


Figure 3: Pre-training then RL training experiments with the DDPT model in interaction with the rule-based simulator. Shaded regions show standard error. Each model is evaluated on 9 different seeds.




Figure 4: Evaluation of the PPO policy trained combined with a SetSUMBT DST model with and without uncertainty features respectively. The policy is trained in an environment that contains 5% user NLG noise to illustrate the impact of uncertainty.

policy trained without these features. [van Niekerk et al. \(2021\)](#) further showed that the policy trained with uncertainty features performs significantly better in conversation with humans than the policy trained without. This is an indication that the policy using uncertainty features can handle ambiguities in conversation better than the policy without uncertainty modeling. To investigate how this policy resolves uncertainty we analyze the action distributions of the policy using the new RL toolkit evaluation tools, which provide new insights into the behaviour of dialogue policy modules. Figure 4 (b) and (c) show that the policy trained using uncertainty features utilizes significantly more re-

quest actions than the policy without these features. This indicates that the policy aims to resolve uncertainty by requesting information from the user. For instance, if the policy recognizes uncertainty regarding the price range a user has requested, it can resolve this through the use of a request. See [van Niekerk et al. \(2021\)](#) for example dialogues with humans where this can be observed.

## G Configuring the Dialogue Pipeline and User Environment

The modules in the dialogue system pipeline and the user environment for interaction are specified

NLU	MultiWOZ 2.1		SGD		Taskmaster	
	ACC	F1	ACC	F1	ACC	F1
T5-NLU	77.8	86.5	45.0	58.6	81.8	73.0
	77.5	86.4	45.2	58.6	81.8	73.0
DST	JGA	Slot F1	JGA	Slot F1	JGA	Slot F1
T5-DST	52.6	91.9	20.1	58.5	48.5	81.1
	53.1	91.9	20.6	60.0	48.6	81.0
SetSUMBT	50.3	90.8	20.0	58.8	24.9	65.5
	50.8	91.0	21.1	59.2	25.3	67.0
NLG	SER ↓	BLEU	SER ↓	BLEU	SER ↓	BLEU
T5-NLG	3.7	35.8	11.9	29.6	2.1	51.5
	3.2	35.6	8.3	29.9	2.0	51.3
SC-GPT	4.8	33.6	9.6	28.2	2.2	47.9
	3.3	33.5	6.8	29.8	1.6	47.3
End2End	Comb.	BLEU	Slot F1	BLEU	Slot F1	BLEU
SOLOIST	67.0	16.8	56.9	11.2	8.5	28.0
	71.4	17.1	69.7	23.1	9.2	29.2
US-DA	ACC	F1	ACC	F1	ACC	F1
TUS	15.0	53.2	10.2	11.6	23.0	23.0
	32.0	62.3	13.8	15.6	22.5	22.0
US-NL	SER ↓	F1	SER ↓	F1	SER ↓	F1
GenTUS	4.2	62.5	8.4	48.8	4.2	43.8
	3.3	49.5	8.0	48.7	3.5	43.9

Table 7: Comparison of independent training and joint training (1st row vs. 2nd row of each model) on 3 datasets. We normalize the service name to the domain name when evaluating NLU and DST on SGD.

using a configuration file as can be seen in Listing 2.

## H Ontology, Data and Database in the Unified Format

As explained in Section 3, a dataset in the unified format consists of an ontology, dialogues and a database. We depict an example for ontology, dialogues and database results for the MultiWOZ 2.1 dataset in the unified format in Listing 3, 4 and 5, respectively.

**Original input:** user: Yes please, for 8 people at 18:30 on thursday.

**Augmented input:** **tm3** user: Yes, please, four for 8:10pm. => [inform] [movie]([num.tickets][four],[time.showing][8:10 pm]) **tm1** user: Yes, 8PM, please. => [inform][restaurant \_reservation]([time.reservation][8PM]) **sgd** user: Yes please, for 3 people on March 8th at 12:30 pm. => [affirm\_intent] [Restaurants\_1]([]);[inform][Restaurants\_1] ([party\_size][3] [date][March 8th],[time][12:30 pm]) **multiwoz21** user: Yes please, for 8 people at 18:30 on thursday.

**Output:** [inform][restaurant]([book day][thursday],[book time][18:30],[book people][8])

Table 8: An example of input augmented by retrieved top-3 samples from other TOD datasets for in-context learning. Dataset names are highlighted.

T5-NLU	MultiWOZ 2.1					
	1%		10%		100%	
T5-small	ACC	F1	ACC	F1	ACC	F1
Baseline	48.1	64.6	68.8	80.6	77.8	86.5
Pre-trained	55.5	70.1	69.8	81.0	77.9	86.5
<i>Data Aug.</i>						
- top1	51.4	66.8	69.0	80.6	77.5	86.5
- top3	51.3	66.4	68.8	80.6	77.0	86.2
<i>In-context</i>						
- top1	44.5	61.3	68.5	80.1	77.7	86.4
- top3	43.6	60.8	68.1	79.8	77.5	86.4
T5-large	ACC	F1	ACC	F1	ACC	F1
Baseline	51.2	67.9	67.8	80.0	76.8	86.4
Pre-trained	56.7	71.7	69.5	81.0	76.8	86.1
<i>Data Aug.</i>						
- top1	49.7	66.8	68.7	80.6	76.9	86.1
- top3	48.5	66.2	68.5	80.5	76.3	85.8
<i>In-context</i>						
- top1	43.4	61.3	69.1	81.0	76.5	85.9
- top3	43.9	63.0	68.7	80.8	76.9	86.2

Table 9: Comparison of different ways to use other TOD datasets: (1) pre-training, (2) retrieving similar samples for data augmentation or (3) in-context learning.

```

{
  "model": {
    "load_path": "from_pretrained",
    "pretrained_load_path": "",
    "use_pretrained_initialisation": false,
    "batchsz": 200,
    "seed": 0,
    "epoch": 100,
    "eval_frequency": 5,
    "process_num": 1,
    "num_eval_dialogues": 20,
    "sys_semantic_to_usr": false
  },
  "vectorizer_sys": {
    "uncertainty_vector_mul": {
      "class_path": "convlab.policy.vector.vector_binary.VectorBinary",
      "ini_params": {
        "use_masking": true,
        "manually_add_entity_names": true,
        "seed": 0
      }
    }
  },
  "nlu_sys": {
    "BertNLU": {
      "class_path": "convlab.nlu.jointBERT.unified_datasets.BERTNLU",
      "ini_params": {
        "mode": "all",
        "config_file": "multiwoz21_all.json",
        "model_file": "https://huggingface.co/ConvLab/bert-base-nlu/resolve/main/bertnlu_unified_multiwoz21_all_context0.zip"
      }
    }
  },
  "dst_sys": {
    "RuleDST": {
      "class_path": "convlab.dst.rule.multiwoz.dst.RuleDST",
      "ini_params": {}
    }
  },
  "sys_nlg": {},
  "nlu_usr": {},
  "dst_usr": {},
  "policy_usr": {
    "GenTUS": {
      "class_path": "convlab.policy.genTUS.stepGenTUS.UserPolicy",
      "ini_params": {
        "model_checkpoint": "convlab/policy/genTUS/unify/experiments/multiwoz21-exp",
        "mode": "language",
        "only_action": false
      }
    }
  },
  "usr_nlg": {}
}

```

Listing 2: Example configuration file for a PPO dialogue policy, binary vectoriser, BertNLU and rule-based DST. The user simulator is GenTUS. We can straightforwardly build different configurations by substituting different modules.

```

{
  "domains": {
    "attraction": {
      "description": "find an attraction",
      "slots": {
        "area": {
          "description": "area to search for attractions",
          "is_categorical": true,
          "possible_values": ["centre", "east", "north", "south", "west"]
        },
        "name": {
          "description": "name of the attraction",
          "is_categorical": false,
          "possible_values": []
        }
      },
      ...
    },
    ...
  },
  "intents": {
    "inform": {"description": "inform the value of a slot"},
    "request": {"description": "ask for the value of a slot"},
    ...
  },
  "state": {
    "attraction": {
      "type": "",
      "name": "",
      "area": ""
    },
    "hotel": {
      "name": "",
      "area": "",
      ...
    },
    ...
  },
  "dialogue_acts": {
    "categorical": [
      "{ 'user': False, 'system': True, 'intent': 'nobook', 'domain': 'hotel', 'slot': 'book day' }",
      "{ 'user': False, 'system': True, 'intent': 'nobook', 'domain': 'restaurant', 'slot': 'book day' }",
      ...
    ],
    "non-categorical": [
      "{ 'user': False, 'system': True, 'intent': 'inform', 'domain': 'attraction', 'slot': 'address' }",
      "{ 'user': False, 'system': True, 'intent': 'inform', 'domain': 'attraction', 'slot': 'choice' }",
      ...
    ],
    "binary": [
      "{ 'user': False, 'system': True, 'intent': 'book', 'domain': 'attraction', 'slot': '' }",
      "{ 'user': False, 'system': True, 'intent': 'book', 'domain': 'hospital', 'slot': '' }",
      ...
    ]
  }
}

```

Listing 3: Example of the unified format ontology for the MultiWOZ 2.1 dataset.



```

[
  {
    "dataset": "multiwoz21",
    "data_split": "train",
    "dialogue_id": "multiwoz21-train-0",
    "original_id": "SNG01856.json",
    "domains": ["hotel", "general"],
    "goal": {
      "description": "You are looking for a place to stay. The hotel should be in the cheap price
range and should be in the type of hotel...",
      "inform": {
        "hotel": {
          "type": "hotel",
          "parking": "yes",
          "price range": "cheap",
          "internet": "yes",
          "book stay": "3|2",
          "book day": "tuesday",
          "book people": "6"
        }
      },
      "request": {
        "hotel": {}
      }
    },
    "turns": [
      {
        "speaker": "user",
        "utterance": "am looking for a place to to stay that has cheap price range it should be in a
type of hotel",
        "utt_idx": 0,
        "dialogue_acts": {
          "categorical": [
            {
              "intent": "inform",
              "domain": "hotel",
              "slot": "price range",
              "value": "cheap"
            }
          ],
          "non-categorical": [
            {
              "intent": "inform",
              "domain": "hotel",
              "slot": "type",
              "value": "hotel",
              "start": 87,
              "end": 92
            }
          ],
          "binary": []
        },
        "state": {
          "attraction": {
            "type": "",
            "name": "",
            "area": ""
          },
          ...
        }
      },
      {
        "speaker": "system",
        "utterance": "Okay, do you have a specific area you want to stay in?",
        "utt_idx": 1,
        "dialogue_acts": {
          "categorical": [],
          "non-categorical": [],
          "binary": [
            {
              "intent": "request",
              "domain": "hotel",
              "slot": "area"
            }
          ]
        },
        "booked": {
          "taxi": [],
          "restaurant": [],
          ...
        }
      },
      ...
    ]
  },
  ...
]

```

Listing 4: Example of the unified format data within the MultiWOZ 2.1 dataset.

```
[
  {
    "address": "124 tenison road",
    "area": "east",
    "internet": "yes",
    "parking": "no",
    "id": "0",
    "name": "a and b guest house",
    "phone": "01223315702",
    "postcode": "cb12dp",
    "price": {
      "double": "70",
      "family": "90",
      "single": "50"
    },
    "pricerange": "moderate",
    "stars": "4",
    "takesbookings": "yes",
    "type": "guesthouse",
    "Ref": "00000000"
  },
  ...
]
```

Listing 5: Example of database query result when searching for a moderately priced hotel in the east from the MultiWOZ unified format database.