

Universal Post-Processing Networks for Joint Optimization of Modules in Task-Oriented Dialogue Systems

Atsumoto Ohashi, Ryuichiro Higashinaka

Graduate School of Informatics, Nagoya University, Japan
 ohashi.atsumoto.c0@s.mail.nagoya-u.ac.jp, higashinaka@i.nagoya-u.ac.jp

Abstract

Post-processing networks (PPNs) are components that modify the outputs of arbitrary modules in task-oriented dialogue systems and are optimized using reinforcement learning (RL) to improve the overall task completion capability of the system. However, previous PPN-based approaches have been limited to handling only a subset of modules within a system, which poses a significant limitation in improving the system performance. In this study, we propose a joint optimization method for post-processing the outputs of all modules using universal post-processing networks (UniPPNs), which are language-model-based networks that can modify the outputs of arbitrary modules in a system as a sequence-transformation task. Moreover, our RL algorithm, which employs a module-level Markov decision process, enables fine-grained value and advantage estimation for each module, thereby stabilizing joint learning for post-processing the outputs of all modules. Through both simulation-based and human evaluation experiments using the MultiWOZ dataset, we demonstrated that UniPPN outperforms conventional PPNs in the task completion capability of task-oriented dialogue systems.

Code — <https://github.com/nu-dialogue/UniPPN>

1 Introduction

Typical task-oriented dialogue systems process a single user input through multiple subtasks to produce a final response. These subtasks include (1) natural language understanding (NLU), which estimates the user’s intent from the input; (2) dialogue state tracking (DST), which accumulates the user’s requests up to the current turn as a dialogue state; (3) dialogue policy (policy), which determines the next action that the system should take as dialogue acts (DAs); and (4) natural language generation (NLG), which converts these DAs into a final system response. Recent research has moved beyond optimizing dedicated modules for each subtask individually. It has explored the use of reinforcement learning (RL) to train several modules based on actual dialogue experiences, thereby optimizing the overall task completion capability of the system (Ni et al. 2022; Kwan et al. 2023).

Recently, a new approach using post-processing networks (PPNs) was proposed to optimize the task completion capability of dialogue systems without directly training the




Figure 1: Diagram of UniPPN. UniPPN modifies the output out_m of Module $_m$ to out_m^+ , which serves as the input for Module $_{m+1}$.

modules (Ohashi and Higashinaka 2022). In this approach, instead of training the modules directly, PPNs, which are components with trainable parameters that modify their outputs, are trained using RL. For example, Ohashi and Higashinaka (2022) implemented PPNs that post-process outputs from NLU, DST, and policy as multi-binary classification tasks using multi-layer perceptrons (referred to as BinPPN), and demonstrated that this optimization improved the overall task completion capability of the systems. In addition, a large language model (LLM)-based generative PPN (GenPPN) was proposed to post-process the natural language output from NLG, and its effectiveness has been demonstrated (Ohashi and Higashinaka 2023).

However, conventional PPN-based methods have two major limitations. First, BinPPN and GenPPN cannot be optimized jointly because of their different model architectures and training algorithms. Although it is possible to post-process the outputs of all modules by combining disjointly trained BinPPNs and GenPPNs, using multiple networks that are not jointly optimized may fail to achieve sufficient performance improvement. The second limitation is the narrow applicability of GenPPN. GenPPN aims to generate utterances that are easily understood by users and relies on a reward function that requires feedback on whether the DAs output from the policy are correctly conveyed to the user. This approach cannot be applied to an end-to-end dialogue system, such as the LLM-powered model proposed by Hudeček and Dusek (2023), which does not explicitly output DAs, imposing significant limitations on GenPPN’s applicability.

In this study, we propose a universal post-processing network (UniPPN) and an optimization method that com-

bines the strengths of two conventional PPNs. UniPPN can jointly optimize the post-processing of outputs from all modules in task-oriented dialogue systems (Figure 1). Our proposed method involves a single-language model-based UniPPN that post-processes the outputs from all modules as a sequence-transformation task. Additionally, we introduce a newly designed module-level Markov decision process (MDP) that extends the standard MDP paradigm and incorporates it into UniPPN’s optimization algorithm. This approach enables fine-grained value and advantage estimation for each module, even with sparse feedback obtained only at the end of multi-turn dialogues, thus ensuring stable joint optimization. Because UniPPN does not require as dense feedback about DAs as GenPPN, it can be applied to a wide range of systems, including end-to-end systems that do not output DAs.

To verify the effectiveness of UniPPN across various dialogue systems, we conducted experiments using dialogue simulations based on the MultiWOZ dataset (Budzianowski et al. 2018). Specifically, we compared the task completion capabilities of dialogue systems equipped with conventional PPNs to those using UniPPN. The results demonstrated that UniPPN significantly outperformed conventional PPNs in terms of task completion capability. Additionally, through dialogue experiments with human users, we demonstrated that dialogue systems using UniPPN outperformed those using conventional PPNs.

2 Related Work

Modularity of Task-Oriented Dialogue Systems In typical pipeline task-oriented dialogue systems, dedicated modules for each subtask, such as NLU, DST, policy, and NLG, have been individually developed and optimized (Zhang et al. 2020). However, in recent years, methods addressing multiple subtasks using a single model have become common (Ni et al. 2022). For example, word-level DST (Wu et al. 2019; Zhao et al. 2022) estimates the dialogue state directly from the dialogue history without requiring user intent estimation using NLU. Similarly, a word-level policy (Lubis et al. 2020; Wang et al. 2020) generates system responses directly from the dialogue state without requiring conversion from DA to system utterances by NLG. Furthermore, end-to-end dialogue systems (Hosseini-Asl et al. 2020; He et al. 2022; Wu et al. 2023), which learn all subtasks using a single model, are becoming popular. Because these end-to-end systems maintain modularity by sequentially executing each subtask, they are often referred to as modularly end-to-end systems (Qin et al. 2023).

Online RL for Task Completion In response generation for task-oriented dialogue systems, it is crucial not only to maximize the probability of reference tokens in corpora but also to maximize task completion capability in actual multi-turn dialogues (Kwan et al. 2023). Some studies (Liu et al. 2018; Tseng et al. 2021) employed online RL frameworks to train dialogue systems based on experiences obtained online from interactions with users. For example, research has focused on optimizing DST (Chen et al. 2022) or policy (Li et al. 2020; Deng et al. 2024) within pipeline systems. Addi-

tionally, Zhao and Eskenazi (2016) demonstrated that jointly optimizing DST and policy with shared parameters outperforms systems in which the DST and policy are trained separately. Our proposed method also utilizes an online RL framework to optimize dialogue systems. However, contrary to previous studies that focused on learning the modules, we concentrate on learning to modify the outputs of these modules.

Post-Processing Networks Methods that train modules via RL cannot be applied to dialogue systems with non-trainable modules, such as rule-based or API-based modules, as expected in real-world scenarios. To address this issue, Ohashi and Higashinaka (2022) proposed optimizing BinPPNs instead of the modules. BinPPNs modify the outputs of NLU, DST, and policy, and are optimized using online RL. Specifically, BinPPNs perform post-processing on the set of slot-value pairs output by each module through binary classification to determine whether to delete (0) or maintain (1) each pair. To handle post-processing for NLG, which outputs natural language rather than a set of slot-value pairs, Ohashi and Higashinaka (2023) introduced GenPPN, which uses LLMs to paraphrase the system utterance output by NLG, to improve task completion by generating system utterances that can be easily understood by users. It is optimized through RL based on feedback regarding whether DAs are correctly conveyed to users.

Our proposed UniPPN can process arbitrary sequences as both input and output. Therefore, contrary to BinPPN, which is limited to binary decisions of deletes/maintenance, it allows for more flexible post-processing. In addition, contrary to GenPPN, which is optimized using detailed feedback on DAs, UniPPN is optimized solely using the task success/failure signal obtained at the end of the dialogue. This makes it applicable to a wide range of systems, including word-level policies and end-to-end systems.

3 Preliminary

The problem of learning capabilities for multi-turn task-oriented dialogue is often formulated as an MDP and optimized through RL. An MDP is defined by tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma)$. Essentially, \mathcal{S} and \mathcal{A} represent all possible dialogue histories and system response sentences, respectively. $\mathcal{P}(s'|s, a)$ represents the transition model, $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ defines the dialogue environment containing the user, and $R(s, a)$ represents the immediate reward function $\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. γ denotes the discount factor. At each turn t , the policy $\mathcal{F} : \mathcal{S} \rightarrow \mathcal{A}$ (i.e., the dialogue system) samples an action (i.e., the system response) $a_t \sim \mathcal{F}(a_t|s_t)$. Until the final state at turn T is reached, the next state $s'_t \sim P(s'_t|s_t, a_t)$ and the immediate reward $r_t = R(s_t, a_t)$ are obtained. The goal of RL is to train \mathcal{F} to maximize the value function V , which is the expected cumulative discounted reward as follows.

$$V^{\mathcal{F}}(s) := \mathbb{E} \left[\sum_{t=0}^T \gamma^t r_t | s_0 = s \right] \quad (1)$$

Numerous studies targeted only a part of \mathcal{F} , such as the policy module, rather than the entire \mathcal{F} .

In complex problems, such as task-oriented dialogues, directly obtaining a policy that maximizes Eq. (1) is challeng-

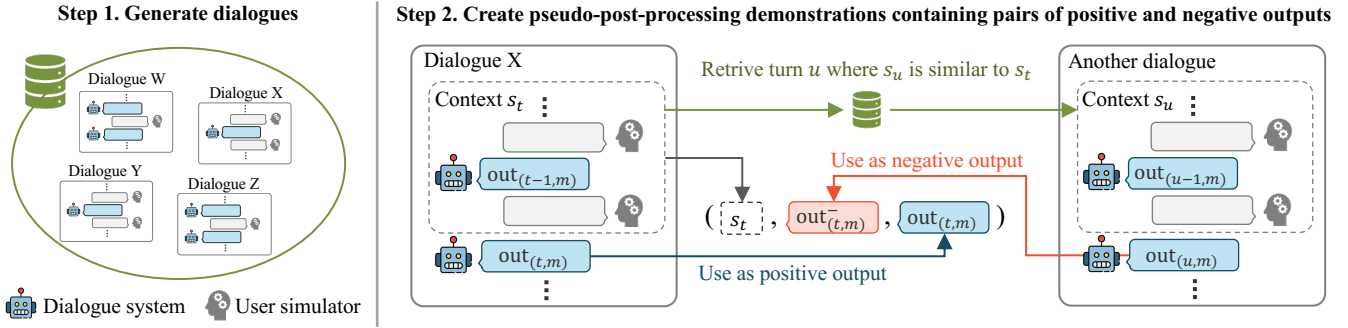


Figure 2: Procedure for creating pseudo-post-processing demonstration data. First, we generate dialogues between the dialogue system and the user simulator. Subsequently, we create pairs of positive and negative outputs, where the output out_t of module m for context s_t at turn t is positive and the output out_u at another turn u is negative (i.e., out_t^-). In imitation learning stage, the reconstruction from out_t^- to out_t is learned as pseudo-post-processing.

ing. One effective method is the policy-gradient-based approach (Sutton et al. 1999), which directly improves the policy network \mathcal{F}_θ parameterized by θ . According to the policy gradient theorem, gradient $\nabla_{\mathcal{F}} J(\theta)$ is expressed as follows:

$$\nabla_{\mathcal{F}} J(\theta) = \mathbb{E} \left[\sum_{t=0}^T \Psi_t \nabla_{\theta} \log \mathcal{F}_\theta(a_t | s_t) \right] \quad (2)$$

The specific definition of Ψ_t varies depending on the implementation of the RL algorithm, such as the sum of the rewards obtained over all turns or advantage estimates (Schulman et al. 2015).

4 Proposed Method

In this section, we explain the problem formulation of our study, the proposed UniPPN, imitation learning (IL) and RL, which together constitute our optimization procedure for UniPPN.

Problem Formulation

Here, we formulate the optimization problem for the dialogue system \mathcal{F} through post-processing. We assume that \mathcal{F} has a modularity consisting of M modules: Module₁, ..., Module_M. At each turn t , each module Module _{m} takes the output $out_{(t,m-1)}$ of the previous Module _{$m-1$} as its input $in_{(t,m)}$, outputting its processing result $out_{(t,m)} \sim \text{Module}_m(in_{(t,m)})$. Some modules may use the dialogue history s_t as additional input. The post-processing network PPN _{m} for Module _{m} modifies $out_{(t,m)}$ and the modified $out_{(t,m)}^+ \sim \text{PPN}_m(s_t, in_{(t,m)}, out_{(t,m)})$ becomes the input for Module _{$m+1$} . For the optimization of \mathcal{F} , we train PPN _{m} instead of Module _{m} .

UniPPN

In our proposed method, the post-processing of the outputs from all M modules is performed by a single network, UniPPN π (Figure 1). Specifically, UniPPN modifies the output of any Module _{m} : $out_{(t,m)}^+ \sim \text{UniPPN}(s_t, in_{(t,m)}, out_{(t,m)}, prefix_m)$. Here, $prefix_m$ is an indicator that specifies that the module to be post-processed is Module _{m} . The

input and output formats of UniPPN are text sequences, with post-processing executed as a sequence-transformation task. For the tokenized sequences $\mathbf{x} = (x_1, \dots, x_k)$ and $\mathbf{y} = (y_1, \dots, y_l)$, representing $(s_t, in_{(t,m)}, out_{(t,m)}, prefix_m)$ and $out_{(t,m)}^+$ respectively, the following conditional probability is modeled:

$$\pi_\theta(\mathbf{y} | \mathbf{x}) = \prod_{i=1}^l \pi_\theta(y_i | \mathbf{x}, \mathbf{y}_{<i}) \quad (3)$$

where π_θ represents a pre-trained language model parameterized by θ . By treating not only the post-processing of natural language, such as the output of NLG modules but also structural data, such as the output of NLU or DST as a sequence-transformation task (Raffel et al. 2020; Liang et al. 2020), UniPPN can uniformly perform post-processing across all modules.

Imitation Learning of Post-Processing

Pre-trained language models are typically trained on web text and may not sufficiently possess the ability to modify the outputs of modules in task-oriented dialogue systems. Therefore, we conduct additional pre-training to teach the model π_θ the formats of input $(s_t, in_{(t,m)}, out_{(t,m)}, prefix_m)$ and output $out_{(t,m)}^+$ through supervised fine-tuning. In the general RL paradigm, IL conducted before online RL uses demonstration data, which consist of the action history of experts, such as humans. However, in our problem setting, demonstration data for post-processing the outputs of each module in the dialogue system \mathcal{F} do not exist. Therefore, we automatically generate post-processing demonstration data and use them for supervised fine-tuning.

Using the procedure shown in Figure 2, we create pseudo-post-processing demonstration data for each Module _{m} . This process involves sampling dialogues by repeating interactions between \mathcal{F} and the environment \mathcal{P} to generate the input-output history $h_{(t,m)} = (s_t, in_{(t,m)}, out_{(t,m)})$ for each Module _{m} at each turn t , resulting in $H_m = \{h_{(1,m)}, \dots, h_{(|H_m|,m)}\}$. We now demonstrate the modification of $out_{(t,m)}$. Here, the label $out_{(t,m)}^+$, which represents the correct modification of $out_{(t,m)}$, cannot be created automat-

ically. Under the assumption that the output $\text{out}_{(t,m)}$ of Module $_m$ is reasonably valid, we consider $\text{out}_{(t,m)}$ to be the target output after post-processing; we use $\text{out}_{(t,m)}^-$, randomly sampled from another turn u (which may be from the same or a different dialogue) as the negative output that should be post-processed. This creates one demonstration instance $d_{(t,m)} = \{(s_t, \text{in}_{(t,m)}, \text{out}_{(t,m)}^-, \text{prefix}_m), \text{out}_{(t,m)}\}$, representing the modification from $\text{out}_{(t,m)}^-$ to $\text{out}_{(t,m)}$. We applied this pseudo-data creation process to all samples in H_m , resulting in the final demonstration dataset $D_m = \{d_{(1,m)}, \dots, d_{(|H_m|,m)}\}$. In the following section, we describe the two techniques used to create D_m .

Sampling Realistic out^- If we sample a turn u that is completely irrelevant to the context of t , it could introduce noise, causing π to potentially learn to ignore $\text{out}_{(t,m)}^-$ rather than to modify it appropriately. To ensure that the mistakes are reasonable, we sample turns with contexts similar to t . Specifically, from the entire history excluding $h_{(t,m)}$ (i.e., $H_m \setminus \{h_{(t,m)}\}$), we extract the top few turns with the highest cosine similarity to the vector representation of the context s_t in $h_{(t,m)}$, and randomly sample $h_{(u,m)}$ from the extracted turns. We use a general-purpose embedding model, such as E5 (Wang et al. 2022) to vectorize the context.

Learning to Copy In post-processing, it is not always necessary to modify the outputs; outputs without issues should be “copied” without modification. To reflect this, during the IL phase, we input the original $\text{out}_{(t,m)}$ into π to ensure that modifications are not always required. In these cases, the target output is only the special token “COPY”. Specifically, demonstrations of such cases are $d_{(t,m)} = \{(s_t, \text{in}_{(t,m)}, \text{out}_{(t,m)}, \text{prefix}_m), \text{COPY}\}$. This approach allows the model to explicitly learn whether post-processing is necessary, while also reducing the generation costs when post-processing is unnecessary. Whether each instance becomes a copy instance is determined randomly using copy ratio $\alpha \in [0, 1]$, which is a hyperparameter.

We update θ based on the maximum likelihood objective using the final dataset $D_{1:M} = [D_1; \dots; D_M]$, which combines the pseudo-post-processing data for all M modules. The optimized parameters in this IL step are denoted by ϕ .

Optimization with Reinforcement Learning

In the RL phase, we install the UniPPN π_ϕ obtained from the IL step into the dialogue system \mathcal{F} . Subsequently, let \mathcal{F} interact repeatedly with the environment \mathcal{P} over multiple turns and update ϕ based on these experiences using a policy-gradient-based approach. In typical task-oriented dialogue systems using online RL, only a single policy network (e.g., a policy module) operates per turn, and it is updated according to Eq. (2). In contrast, our study involves a policy π that acts M times per turn, and outputs the system response as action a . Although each of the M actions should have different gradients based on their individual advantages, Eq. (2) treats them as having the same contributions. This can result in coarse rewards and learning instability.

Therefore, we extend the standard MDP described in Section 3 and introduce a module-level MDP, where the unit of time step is the “post-processing of one module by π ” rather than the “one turn response by \mathcal{F} ”. Specifically, the value function to be maximized and policy gradient of π_ϕ are as follows:

$$V^\pi(\mathbf{x}) := \mathbb{E} \left[\sum_{t=0}^T \sum_{m=1}^M \gamma^{(t+1)(m-1)} r_{(t,m)} | \mathbf{x}_{(0,1)} = \mathbf{x} \right] \quad (4)$$

$$\nabla_\pi J(\phi) = \mathbb{E} \left[\sum_{t=0}^T \sum_{m=1}^M \Psi_{(t,m)} \nabla_\phi \log \pi_\phi(\mathbf{y}_{(t,m)} | \mathbf{x}_{(t,m)}) \right] \quad (5)$$

Here, $r_{(t,m)}$ represents the immediate reward for post-processing the output of Module $_m$ at turn t . As in previous studies using online RL (Hou et al. 2021), a small negative fixed value is assigned continuously until the end of the dialogue. $\mathbf{x}_{(t,m)}$ and $\mathbf{y}_{(t,m)}$ are the tokenized sequences of the input text $(s_t, \text{in}_{(t,m)}, \text{out}_{(t,m)}, \text{prefix}_m)$ and output text of UniPPN, respectively. Eq. (5) shows that the gradients can be computed in M gradient accumulation steps. Note that, in Eq. (4), the number of calculations for the value function in the module-level MDP is $T \times M$, resulting in a possible exponential increase in the computational cost. However, because M in a typical task-oriented dialogue system is four at most, this is not a problem in practice.

To implement $\Psi_{(t,m)}$, we adopt a generalized advantage estimation (Schulman et al. 2015). Specifically, we compute the advantage estimate $\hat{A}_{(t,m)}$ based on the value $V_\psi(\mathbf{x}_{(t,m)})$ of $\mathbf{x}_{(t,m)}$ estimated using another language model V_ψ parameterized by ψ as a critic network:

$$\begin{aligned} \hat{A}_{(t,m)} &= \delta_{(t,m)} + \gamma \lambda \hat{A}_{(t,m)'}, \\ \delta_{(t,m)} &= r_{(t,m)} + \gamma V_\psi(\mathbf{x}_{(t,m)'}) - V_\psi(\mathbf{x}_{(t,m)}) \end{aligned} \quad (6)$$

$$(t, m)' = \begin{cases} (t, m + 1) & \text{if } m < M \\ (t + 1, 1) & \text{if } m = M \end{cases} \quad (7)$$

Here, $\delta_{(t,m)}$ represents the TD residual, and the hyperparameter $\lambda \in [0, 1]$ controls the trade-off between utilizing actual long-term rewards and the estimated values. Because V_ψ estimates the state value for each Module $_m$ at each turn t , fine-grained advantage estimation according to the contribution of each module is possible even in settings with sparse rewards across multi-turn dialogues. An advantage of this algorithm is that it does not require a high-cost manual reward design for each module, as required in previous studies. V_ψ is trained to minimize the mean squared error with respect to the cumulative reward and π_ϕ is optimized using a clipped surrogate objective with proximal policy optimization (PPO) (Schulman et al. 2017). For a detailed implementation of the RL algorithm, refer to Appendix A.

5 Experiments

In this evaluation experiment, we demonstrate that joint optimization using UniPPN is more effective than disjoint optimization combining conventional BinPPN and GenPPN for post-processing outputs from all modules to improve task-oriented dialogue systems.

Experimental Setup

We conducted evaluation experiments using the MultiWOZ (Budzianowski et al. 2018) dataset, which contains a multi-domain task-oriented dialogue on travel information between customers and clerks. We applied UniPPN to various dialogue systems developed for MultiWOZ and assessed their task completion capabilities. For the user simulation, we used the agenda-based user simulator (Schatzmann et al. 2007) provided by ConvLab-2 (Zhu et al. 2020), which is an evaluation toolkit for task-oriented dialogue systems.

The dialogue systems used in our experiments included both pipeline and end-to-end systems. For the modules constituting the pipeline systems, we selected relatively recently proposed models that ranked high on the MultiWOZ benchmark¹ and had publicly available implementations. The models adopted for each module in the pipeline system are as follows:

NLU BERT NLU (Chen, Zhuo, and Wang 2019), a classification model based on BERT (Devlin et al. 2019).

DST Rule-based DST and D3ST (Zhao et al. 2022), a state-of-the-art word-level DST based on T5 (Raffel et al. 2020).

Policy Rule-based policy and PPO policy fine-tuned using PPO (Schulman et al. 2017). We also used LAVA (Lubis et al. 2020), a word-level policy.

NLG Template-based NLG and SC-GPT (Peng et al. 2020), which is based on GPT-2 (Radford et al. 2019).

For end-to-end systems, we adopted two representative models: PPTOD (Su et al. 2022) and an LLM-based model (Hudeček and Dusek 2023). PPTOD is a T5-based dialogue model that is fine-tuned using MultiWOZ. The LLM-based model performs word-level DST and a word-level policy based on in-context learning with few-shot examples retrieved from MultiWOZ. For the LLM, we used GPT-4o mini provided by OpenAI’s API.²

Evaluation Metrics

In the evaluation, each dialogue system interacted with the user simulator 1,024 times, and each of the 1,024 different user goals for testing was set in each dialogue. We reported the average score of 1,024 dialogues as the final score.

As evaluation metrics, we used the average *turns* across all dialogues, which represents the number of turns required to achieve a task, with lower values indicating better efficiency. Each turn comprised a pair of a user utterance and the corresponding system response. We also used *Inform Recall/Precision/F1*. These metrics assess whether the system responds adequately to the information requested by the user during a dialogue. In addition, we used the *Goal Match Rate* to assess whether the conditions of the entity (specific facilities, such as hotels) presented by the system matched the user’s goal. Similarly, we also assessed the conditions of the entity booked by the system using the *Book Rate*. We set the maximum number of turns in one dialogue to 20. A task was

only considered *Success* if the Inform Recall, Match Rate, and Book Rate all reached 1.0 within 20 turns.

UniPPN Implementation

We used a medium-sized GPT-2 (Radford et al. 2019) with 355M parameters as the backbone model for UniPPN. We chose this parameter size because of its superior balance between the computational cost and performance, which was confirmed through preliminary experiments.

Imitation Learning To construct the post-processing demonstration data $D_{1:M}$ for each dialogue system, we sampled 10,000 turns of interaction between the dialogue system and user simulator.³ To sample turns with similar contexts in the construction of $D_{1:M}$, we adopted GTE-base (Li et al. 2023) as the embedding model and used the latest three utterances as the context. In addition, we set the copy ratio α to 0.1 throughout the experiment because, in preliminary experiments, we examined 9 levels from 0.1 to 0.9 for α and found that 0.1 yielded the highest reward.

Reinforcement Learning As an approximator for the value function, we used GPT-2 of 124M parameters, with an additional feedforward network outputting a scalar value. For the reward function $R(t, m)$, we set a small negative value of $R(t, m) = -0.1$ until the end of the dialogue and $R(T, M) = 2 \times M$ for the final step if the task was achieved. We trained for 200 iterations, and in each iteration, we sampled 1,024 turns as training data. We used the checkpoints from the final iteration for testing.

For detailed implementations and hyperparameters of the learning process, please refer to Appendix B.

Baselines

As baselines for this experiment, we used two methods: the original dialogue system without post-processing and a method that post-processes the outputs of all modules by combining conventional BinPPN and GenPPN (called BinPPN&GenPPN). Because BinPPN and GenPPN cannot be trained jointly, RL was used to train the two types of PPNs (BinPPN and GenPPN) separately. Specifically, we used RL to optimize the post-processing of the three modules (NLU, DST, and policy) using BinPPN. Thereafter, while installing these three BinPPNs in the system and freezing their parameters, we attached GenPPN to NLG and trained it again using RL. The BinPPN was optimized first because NLU, DST, and policy, whose output BinPPN post-processes, precede NLG, whose output GenPPN post-processes; therefore, this order of optimization is considered appropriate. Note that, contrary to UniPPN, BinPPN&GenPPN require two RL phases. For the implementation and hyperparameters of BinPPN and GenPPN, we used those published and reported in previous studies (Ohashi and Higashinaka 2022, 2023). However, for the backbone model of GenPPN, we used Llama 3.1

³10,000 turns equal approximately 1,000 dialogues, although the number of turns per dialogue depends on various factors, such as the system’s capabilities.

¹<https://github.com/budzianowski/multiwoz>

²<https://platform.openai.com>

System	Module combination				Success Rate	Inform			Book Rate	Match Rate	Turns ↓
	NLU	DST	Policy	NLG		Recall	Prec.	F1			
<i>Pipeline System</i>											
SYS _{RULE}	BERT	Rule	Rule	Template	83.69	93.72	81.38	85.08	91.19	91.63	5.92
+BinPPN&GenPPN	✓	✓	✓	✓	85.25	93.96	83.53	86.54	91.18	92.68	5.67
+UniPPN	✓	✓	✓	✓	90.62	96.46	82.86	87.19	97.13	94.81	5.77
SYS _{D3ST}	–	D3ST ^{word}	Rule	Template	58.50	72.41	67.36	66.62	62.22	77.12	6.78
+BinPPN&GenPPN	–	✓	✓	✓	67.58	82.24	68.02	71.89	75.31	81.59	6.36
+UniPPN	–	✓	✓	✓	85.06	93.63	71.82	78.56	89.25	92.15	5.78
SYS _{PPO}	BERT	Rule	PPO	Template	69.24	86.90	66.99	72.90	80.74	83.79	8.55
+BinPPN&GenPPN	✓	✓	✓	✓	70.61	86.59	66.08	72.46	80.71	84.42	8.46
+UniPPN	✓	✓	✓	✓	89.36	96.72	68.37	77.68	94.73	94.82	5.24
SYS _{SCGPT}	BERT	Rule	Rule	SC-GPT	75.29	93.73	79.49	83.92	68.53	92.06	6.12
+BinPPN&GenPPN	✓	✓	✓	✓	86.72	94.66	82.79	86.32	92.53	92.84	5.68
+UniPPN	✓	✓	✓	✓	90.14	97.19	84.14	88.38	95.89	95.69	6.10
SYS _{LAVA}	BERT	Rule	LAVA ^{word}	–	64.36	82.87	55.74	64.04	72.96	80.34	10.24
+UniPPN	✓	✓	✓	–	79.39	98.10	64.41	75.21	88.66	89.91	5.84
<i>End-to-End System</i>											
SYS _{PPTOD}	–	PPTOD ^{word}	PPTOD ^{word}	–	61.04	82.19	75.12	76.07	50.80	83.61	8.15
+UniPPN	–	✓	✓	–	80.37	92.71	75.83	81.00	83.08	90.01	6.97
SYS _{LLM} (GPT-4o mini)	–	LLM ^{word}	LLM ^{word}	–	62.30	77.66	58.64	62.85	62.37	79.26	7.66
+UniPPN	–	✓	✓	–	88.28	94.80	65.01	74.31	90.81	93.57	4.66

Table 1: Test results for each dialogue system and when post-processing is applied to all of their modules using BinPPN and GenPPN (+BinPPN&GenPPN) or using UniPPN (+UniPPN). The ✓ under each module in the row indicates whether PPN is applied to that module. Note that UniPPNs applied to the same system are the same network. The superscript “word” for DST and policy indicates that they are word-level DST and word-level policy, respectively.

8B (Dubey et al. 2024) instead of Llama 7B (Touvron et al. 2023) as described in the previous study.

Main Results

Table 1 shows the test results when post-processing was applied to all modules of each system. We evaluated two cases for the application of post-processing to pipeline systems consisting of multiple modules: +BinPPN&GenPPN and +UniPPN. Because BinPPN&GenPPN cannot be applied to SYS_{LAVA} or end-to-end systems that do not output DAs, only UniPPN was applied to these systems.

In the pipeline systems, +UniPPN significantly outperformed +BinPPN&GenPPN. For example, even for systems, such as SYS_{RULE} and SYS_{PPO}, where performance improvement by +BinPPN&GenPPN was limited, UniPPN improved the performance. In particular, it is noteworthy that UniPPN enhanced SYS_{RULE} to 90.62 points, considering that SYS_{RULE} comprises high-performance modules that are carefully crafted by hand-written rules, and its original success rate is high at approximately 84 points. Furthermore, we applied UniPPN to the system, including the word-level policy and end-to-end systems, to which conventional BinPPN or GenPPN could not be applied. UniPPN significantly improved the original performance of SYS_{LAVA}, SYS_{PPTOD}, and SYS_{LLM} for all evaluation metrics.

From these results, we demonstrated that the joint optimization of post-processing the outputs from all modules

using UniPPN is more effective than the disjoint optimization of multiple PPNs. Furthermore, its efficacy was demonstrated, and it was shown that it can be applied to any dialogue system, regardless of the trainability of each module, including rule- or API-based modules. We believe that UniPPN is more practical considering the high training cost of BinPPN&GenPPN and the complexity of installing multiple networks in the systems.

Comparison with either BinPPN or GenPPN

To clarify the fundamental performance difference between UniPPN and BinPPN, we compared the performance when either BinPPN or UniPPN was applied to NLU, DST, and policy, which BinPPN can be applied, for each system. Table 2 presents the results. We can observe that for most systems, the performance improvement by UniPPN exceeds that of BinPPN. This performance difference may be due to limited post-processing capabilities of BinPPN, which is limited to basic binary operations, leaving little room for improvement, whereas UniPPN can flexibly generate various types of information. Similarly, we compared the performance when either GenPPN or UniPPN was applied to the NLG of each system. Table 3 shows the results. For the turn metric, GenPPN consistently outperformed UniPPN. This could be because GenPPN learns to generate responses, such that DAs are easily understood by the user simulator, thereby reducing the increase in turns caused by users asking back.

System	PPN	Success	Inf. F1	Book	Match	Turns ↓
SYS _{RULE}	Bin	83.40	86.07	91.13	91.47	6.04
	Uni	87.11	83.81	96.08	93.13	5.80
SYS _{D3ST}	Bin	62.89	70.78	70.71	79.02	7.21
	Uni	77.93	76.97	79.50	88.66	5.83
SYS _{PP0}	Bin	69.82	72.80	80.78	83.64	8.63
	Uni	85.84	76.37	94.05	92.50	5.40
SYS _{SCGPT}	Bin	73.44	83.60	63.06	91.81	6.08
	Uni	77.44	85.06	69.34	93.62	6.10
SYS _{LAVA}	Bin	64.06	69.56	80.99	79.98	8.72
	Uni	65.72	67.55	76.48	80.83	9.45

Table 2: Test results when applying either BinPPN or UniPPN to the three modules of NLU, DST, and policy in each system. For SYS_{LAVA}, which uses word-level policy, BinPPN cannot be applied, therefore, either BinPPN or UniPPN is applied only to NLU and DST.

System	PPN	Success	Inf. F1	Book	Match	Turns ↓
SYS _{RULE}	Gen	85.06	85.08	90.90	92.35	5.61
	Uni	85.06	85.08	91.58	92.29	5.81
SYS _{SCGPT}	Gen	84.08	85.40	91.62	91.65	5.67
	Uni	84.77	81.66	90.50	92.56	6.15

Table 3: Test results when applying either GenPPN or UniPPN only to the NLG of each system.

However, for other metrics, such as success rate, there was no significant difference between the two methods. Considering that the training algorithm of GenPPN requires the internal DAs of the system and feedback on whether the user understood those DAs, UniPPN, which only requires the final dialogue outcome as a reward, is promising.

Based on these results, we demonstrated that UniPPN addresses the multiple challenges of conventional BinPPN and GenPPN, resulting in improved task completion capability, reduced computational cost, and broader applicability.

6 Human Evaluation

We verified whether the performance improvement of dialogue systems by UniPPN is also effective for human users. Specifically, for SYS_{PP0} in Table 1, we had human users interact with three types of systems (i.e., the original system, +BinPPN&GenPPN, and +UniPPN) and evaluated their performance. We chose SYS_{PP0} because it achieved the lowest task success rate among the systems containing all four modules. We believe that the impact of post-processing on the performance is most apparent in this case. We recruited more than 40 workers for each system on Amazon Mechanical Turk (AMT) and had them interact with one of the three systems to achieve dialogue goals randomly created for each dialogue. After the dialogue, the workers were requested to subjectively evaluate the system’s language understanding (LU), response appropriateness (RA), and overall satisfaction (OS) with the dialogue on a 5-point Likert scale. Eth-

System	<i>N</i>	Success	Turns ↓	LU	RA	OS
SYS _{PP0}	43	46.51	8.79	2.70	2.70	2.63
+BinPPN&GenPPN	40	42.50	9.80	2.70	2.85	2.38
+UniPPN	48	54.17	8.06 ⁺	2.81	2.90	2.56

Table 4: Results of human evaluation. *N* indicates the number of subjects who conversed with each system. LU, RA, and OS indicate evaluations of the system’s language understanding, response appropriateness, and overall satisfaction, respectively. + indicates that there was a significant tendency with $p < 0.1$ according to the Mann-Whitney U test in the difference between the scores of +BinPPN&GenPPN and +UniPPN.

ical approval was obtained from our institution before the experiment. For detailed experimental settings, please refer to Appendix C.

Table 4 lists the results of the human evaluation metrics for the task completion and subjective assessments. For the evaluation metrics related to task completion capability, such as success rate and number of turns, +UniPPN outperformed both the original SYS_{PP0} and +BinPPN&GenPPN. Notably, the difference in the number of turns showed a statistically significant tendency, highlighting the effectiveness of UniPPN, which can jointly learn to post-process the outputs of all modules, compared with disjointly trained conventional PPNs. By contrast, +BinPPN&GenPPN performed worse than the original SYS_{PP0} for most metrics. This may be because the post-processing for each module is trained disjointly, preventing coordination between PPNs and deteriorating the overall system performance. Regarding the subjective evaluation metrics of LU, RA, and OS, there were no significant differences between SYS_{PP0} and +UniPPN. This was expected, considering that UniPPN’s reward signals were only related to task completion and did not include signals related to dialogue satisfaction.

7 Conclusion

In this study, we proposed UniPPN, a method that jointly learns the post-processing of outputs from all modules in task-oriented dialogue systems. Using simulation experiments based on the MultiWOZ dataset, we applied UniPPN to various pipeline systems with recent high-performance modules and end-to-end systems, including a GPT-4o mini-powered system. Our results confirm that UniPPN significantly outperforms conventional PPN-based methods in terms of task completion performance of dialogue systems. Furthermore, human evaluation experiments demonstrated that UniPPN, optimized in a simulation environment, is effective in real-world scenarios.

In future studies, we aim to reduce the overall training cost of UniPPN. This involves reducing the number of dialogue experiences required for learning convergence and optimizing the model size for efficiency. In addition, we plan to extend beyond text dialogue to support the post-processing of outputs from modules in multimodal dialogue systems.

Acknowledgments

This work was supported by JST Moonshot R&D, Grant number JPMJMS2011. We used the computational resources of the supercomputer “Flow” at the Information Technology Center, Nagoya University.

References

- Budzianowski, P.; Wen, T.-H.; Tseng, B.-H.; Casanueva, I.; Ultes, S.; Ramadan, O.; and Gašić, M. 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*, 5016–5026.
- Chen, Q.; Zhuo, Z.; and Wang, W. 2019. BERT for Joint Intent Classification and Slot Filling. *arXiv preprint arXiv:1902.10909*.
- Chen, Z.; Chen, L.; Zhou, X.; and Yu, K. 2022. Deep Reinforcement Learning for On-line Dialogue State Tracking. In *Proceedings of the 17th National Conference on Man-Machine Speech Communication*, 278–292.
- Deng, Y.; Zhang, W.; Lam, W.; Ng, S.-K.; and Chua, T.-S. 2024. Plug-and-Play Policy Planner for Large Language Model Powered Dialogue Agents. In *Proceedings of the Twelfth International Conference on Learning Representations*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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*, 4171–4186.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.
- He, W.; Dai, Y.; Zheng, Y.; Wu, Y.; Cao, Z.; Liu, D.; Jiang, P.; Yang, M.; Huang, F.; Si, L.; Sun, J.; and Li, Y. 2022. GALAXY: A Generative Pre-trained Model for Task-Oriented Dialog with Semi-supervised Learning and Explicit Policy Injection. In *Proceedings of 36th AAAI Conference on Artificial Intelligence*, 10, 10749–10757.
- Hosseini-Asl, E.; McCann, B.; Wu, C.-S.; Yavuz, S.; and Socher, R. 2020. A Simple Language Model for Task-Oriented Dialogue. In *Proceedings of Advances in Neural Information Processing Systems*, 20179–20191.
- Hou, Z.; Liu, B.; Zhao, R.; Ou, Z.; Liu, Y.; Chen, X.; and Zheng, Y. 2021. Imperfect also Deserves Reward: Multi-Level and Sequential Reward Modeling for Better Dialog Management. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2993–3001.
- Hudeček, V.; and Dusek, O. 2023. Are Large Language Models All You Need for Task-Oriented Dialogue? In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, 216–228.
- Kwan, W.-C.; Wang, H.-R.; Wang, H.-M.; and Wong, K.-F. 2023. A survey on recent advances and challenges in reinforcement learning methods for task-oriented dialogue policy learning. *Machine Intelligence Research*, 20(3): 318–334.
- Li, Z.; Lee, S.; Peng, B.; Li, J.; Kiseleva, J.; de Rijke, M.; Shayandeh, S.; and Gao, J. 2020. Guided Dialogue Policy Learning without Adversarial Learning in the Loop. In *Findings of the 2019 Conference on Empirical Methods in Natural Language Processing*, 2308–2317.
- Li, Z.; Zhang, X.; Zhang, Y.; Long, D.; Xie, P.; and Zhang, M. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*.
- Liang, W.; Tian, Y.; Chen, C.; and Yu, Z. 2020. MOSS: End-to-End Dialog System Framework with Modular Supervision. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 8327–8335.
- Liu, B.; Tür, G.; Hakkani-Tür, D.; Shah, P.; and Heck, L. 2018. Dialogue Learning with Human Teaching and Feedback in End-to-End Trainable Task-Oriented Dialogue Systems. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2060–2069.
- Lubis, N.; Geishausser, C.; Heck, M.; Lin, H.-c.; Moresi, M.; van Niekerk, C.; and Gasic, M. 2020. LAVA: Latent Action Spaces via Variational Auto-encoding for Dialogue Policy Optimization. In *Proceedings of the 28th International Conference on Computational Linguistics*, 465–479.
- Ni, J.; Young, T.; Pandelea, V.; Xue, F.; and Cambria, E. 2022. Recent advances in deep learning based dialogue systems: A systematic survey. *Artificial intelligence review*, 1–101.
- Ohashi, A.; and Higashinaka, R. 2022. Post-processing Networks: Method for Optimizing Pipeline Task-oriented Dialogue Systems using Reinforcement Learning. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 1–13.
- Ohashi, A.; and Higashinaka, R. 2023. Enhancing Task-oriented Dialogue Systems with Generative Post-processing Network. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 3815–3828.
- Peng, B.; Zhu, C.; Li, C.; Li, X.; Li, J.; Zeng, M.; and Gao, J. 2020. Few-shot Natural Language Generation for Task-Oriented Dialog. In *Findings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 172–182.
- Qin, L.; Pan, W.; Chen, Q.; Liao, L.; Yu, Z.; Zhang, Y.; Che, W.; and Li, M. 2023. End-to-end Task-oriented Dialogue: A Survey of Tasks, Methods, and Future Directions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 5925–5941.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring

- the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 1–67.
- Schatzmann, J.; Thomson, B.; Weilhammer, K.; Ye, H.; and Young, S. 2007. Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System. In *Proceedings of Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics*, 149–152.
- Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2015. High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv preprint arXiv:1506.02438*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*.
- Su, Y.; Shu, L.; Mansimov, E.; Gupta, A.; Cai, D.; Lai, Y.-A.; and Zhang, Y. 2022. Multi-Task Pre-Training for Plug-and-Play Task-Oriented Dialogue System. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 4661–4676.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 1999. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Proceedings of Advances in Neural Information Processing Systems*, 1057–1063.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Tseng, B.-H.; Dai, Y.; Kreyssig, F.; and Byrne, B. 2021. Transferable Dialogue Systems and User Simulators. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 152–166.
- Wang, K.; Tian, J.; Wang, R.; Quan, X.; and Yu, J. 2020. Multi-Domain Dialogue Acts and Response Co-Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7125–7134.
- Wang, L.; Yang, N.; Huang, X.; Jiao, B.; Yang, L.; Jiang, D.; Majumder, R.; and Wei, F. 2022. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*.
- Wu, C.-S.; Madotto, A.; Hosseini-Asl, E.; Xiong, C.; Socher, R.; and Fung, P. 2019. Transferable Multi-Domain State Generator for Task-Oriented Dialogue Systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 808–819.
- Wu, Q.; Gung, J.; Shu, R.; and Zhang, Y. 2023. DiactTOD: Learning Generalizable Latent Dialogue Acts for Controllable Task-Oriented Dialogue Systems. In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 255–267.
- Zhang, Z.; Takanobu, R.; Zhu, Q.; Huang, M.; and Zhu, X. 2020. Recent advances and challenges in task-oriented dialog systems. *Science China Technological Sciences*, 63(10): 1–17.
- Zhao, J.; Gupta, R.; Cao, Y.; Yu, D.; Wang, M.; Lee, H.; Rastogi, A.; Shafran, I.; and Wu, Y. 2022. Description-driven task-oriented dialog modeling. *arXiv preprint arXiv:2201.08904*.
- Zhao, T.; and Eskenazi, M. 2016. Towards End-to-End Learning for Dialog State Tracking and Management using Deep Reinforcement Learning. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 1–10.
- Zhu, Q.; Zhang, Z.; Fang, Y.; Li, X.; Takanobu, R.; Li, J.; Peng, B.; Gao, J.; Zhu, X.; and Huang, M. 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*, 142–149.