

End-to-End Conversational Search for Online Shopping with Utterance Transfer

Liqiang Xiao^{1*}, Jun Ma², Xin Luna Dong², Pascual Martinez-Gomez²,
Nasser Zalmout², Chenwei Zhang², Tong Zhao², Hao He^{1,3†}, Yaohui Jin^{1,3†}

¹MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

²Amazon.com, Inc.

³State Key Lab of Advanced Optical Communication System and Network,
Shanghai Jiao Tong University

xiaoliqiang@sjtu.edu.cn, {junmaa, lunadong, gomepasc, nzalmout,
cwzhang, zhaoton}@amazon.com, {hehao, jinyh}@sjtu.edu.cn

Abstract

Successful conversational search systems can present natural, adaptive and interactive shopping experience for online shopping customers. However, building such systems from scratch faces real word challenges from both imperfect product schema/knowledge and lack of training dialog data. In this work we first propose **ConvSearch**, an end-to-end conversational search system that deeply combines the dialog system with search. It leverages the text profile to retrieve products, which is more robust against imperfect product schema/knowledge compared with using product attributes alone. We then address the lack of data challenges by proposing an utterance transfer approach that generates dialogue utterances by using existing dialog from other domains, and leveraging the search behavior data from e-commerce retailer. With utterance transfer, we introduce a new conversational search dataset for online shopping. Experiments show that our utterance transfer method can significantly improve the availability of training dialogue data without crowd-sourcing, and the conversational search system significantly outperformed the best tested baseline.

1 Introduction

Search systems play significant roles in today’s online shopping experience. In conventional e-commerce search systems, user interacts with the system through typing of keywords, followed by product clicks or keywords modifications, depending on whether returned product list matches with user expectation. The recent success of intelligent assistants such as Alexa, Google Now, and Siri enables user to interact with search systems using natural language. For online shopping in particular,

it becomes alluring that users can navigate through products with conversations like traditional in-store shopping, guided by a knowledgeable yet thoughtful virtual shopping assistant.

However, building a successful conversational search system for online shopping faces at least two real world challenges. The first challenge is the *imperfect product attribute schema and product knowledge*. While this challenge applies also to traditional search systems, it is more problematic for conversational search because the later depends on product attributes to link lengthy multi-turn utterances (in contrast to short queries in conventional search) with products. Most previous conversational shopping search work (Li et al., 2018a; Bi et al., 2019; Yan et al., 2017) looks for the target product through direct attribute matching, assuming availability of complete product knowledge in structured form. In practice, this assumption rarely holds, and systems designed with this assumption will suffer from product recall losses.

The second challenge is the *lack of in-domain dialog dataset* for model training. Constructing a large-scale dialog dataset by crowd-sourcing from scratch is inefficient. Popular approaches include Machines-Talking-To-Machines (M2M) (Shah et al., 2018), which generates outlines of dialogs by self-play between two machines, and Wizard-of-Oz (WoZ) (Kelley, 1984), which collects data through virtual conversations between annotators. Note that both approaches require manually written utterances. In addition, a line of other work (Lei et al., 2020; Luo et al., 2020; Bi et al., 2019) constructs conversations from the review datasets such as Amazon Product Data (McAuley et al., 2015) and LastFM¹, whereas usage of these

*Work performed during internship at Amazon.

† Corresponding author

¹<https://grouplens.org/datasets/hetrec-2011/>

datasets is limited to sub-tasks (e.g., dialog policy) due to the absence of user utterances. Saha et al. (2018) collected a dialog dataset for fashion product shopping. However, the described method requires dozens of domain experts to manually create the dialog, and the dataset can hardly be generalized beyond fashion shopping given the lack of utterance annotations.

To address the first challenge of imperfect attribute schema and product knowledge, we propose **ConvSearch**, an end-to-end conversational search system that deeply combines the dialog and search system to improve the search performance. In particular, the Product Search module leverages both structured product attributes and unstructured product text (e.g. profile), where the product text may contain phrases matching with utterances when schema is incomplete or when a product attribute value is missing. Putting together, our system has the advantage of both reduced error accumulation along individual modules, and enhanced robustness against product schema/knowledge gaps.

To address the second challenge of lacking in-domain dialog dataset, we propose a jump-start dialog generation method **M2M-UT** which 1) generates utterance from existing dialogues of similar domains (e.g., movie ticketing (Li et al., 2017)), and 2) builds dialog outlines from e-commerce search behavior data, and fills them with the generated utterances. The proposed approach significantly reduces manual effort in data construction, and as a result we introduce a new conversational shopping search dataset **CSD-UT** with 942K utterances. Note that although the dialogue dataset construction focuses on shopping, the approach described here can be adapted for other task-oriented conversations as well, which we will leave it to future work. Our contributions are summarized as follows:

- We proposed an end-to-end conversational search system which deeply combines dialog with search, and leverages both structured product attributes and unstructured text in product search to compensate for incomplete product schema/knowledge.²
- We proposed a new dialog dataset construction approach, which can transfer utterances from dialogs of similar domains and build

²The prototype development, evaluation, and data set presented in this paper are independent from any existing commercialized chatbot system.

dialogues from user behavior records. Using this new approach which significantly reduced manual work compared with existing approaches, we introduced a new conversational search dataset for online shopping.

- Extensive experiments show that our system outperforms evaluated competitors for success rate (SR@5) score.

2 Related Work

Conversational Search System Conversational search task aims to understand user’s search intents through multi-round conversational interactions, and return user the desired search item. Due to lack of annotated dialog utterances in particular for conversational search tasks, previous work either adopted rule-based utterance parsing or focused only on dialog policy. Yan et al. (2017) proposed a rule-based approach to cold-start online shopping dialog systems utilizing user search logs and intent phrases collected from community sites. In another line of work, Luo et al. (2020) and Zhang et al. (2018) utilized Amazon review dataset, Lei et al. (2020) and Li et al. (2018a) revised user reviews from Yelp³ and LastFM⁴, all of which focused on the conversation policy without utterance understanding. As a comparison, in this paper we focused an end-to-end conversational search system, which fuses both utterance understanding and product search together through multi-task learning.

Constructing Dialog Dataset for Online Shopping Rastogi et al. (2020) proposed a crowd-sourcing version of Wizard-of-Oz (WOZ) paradigm for collecting domain-specific corpora. In this system, users and wizards were given a predefined task to complete (e.g. find a Chinese restaurant in the North). To avoid the distracting latency, users and wizards were asked to contribute just a single turn for each dialogue. Saha et al. (2018) built a multi-mode dialog system for fashion with experts and in-house labors. They crawled 1 million fashion items from the web, hand-crafted taxonomy for items, identified the set of fashion attributes, and employed experts to write dialogs. The described methods were highly labor-consuming, and the published dataset did not contain attribute annotation on utterances, making

³<https://www.yelp.com/dataset/>

⁴<https://grouplens.org/datasets/hetrec-2011/>

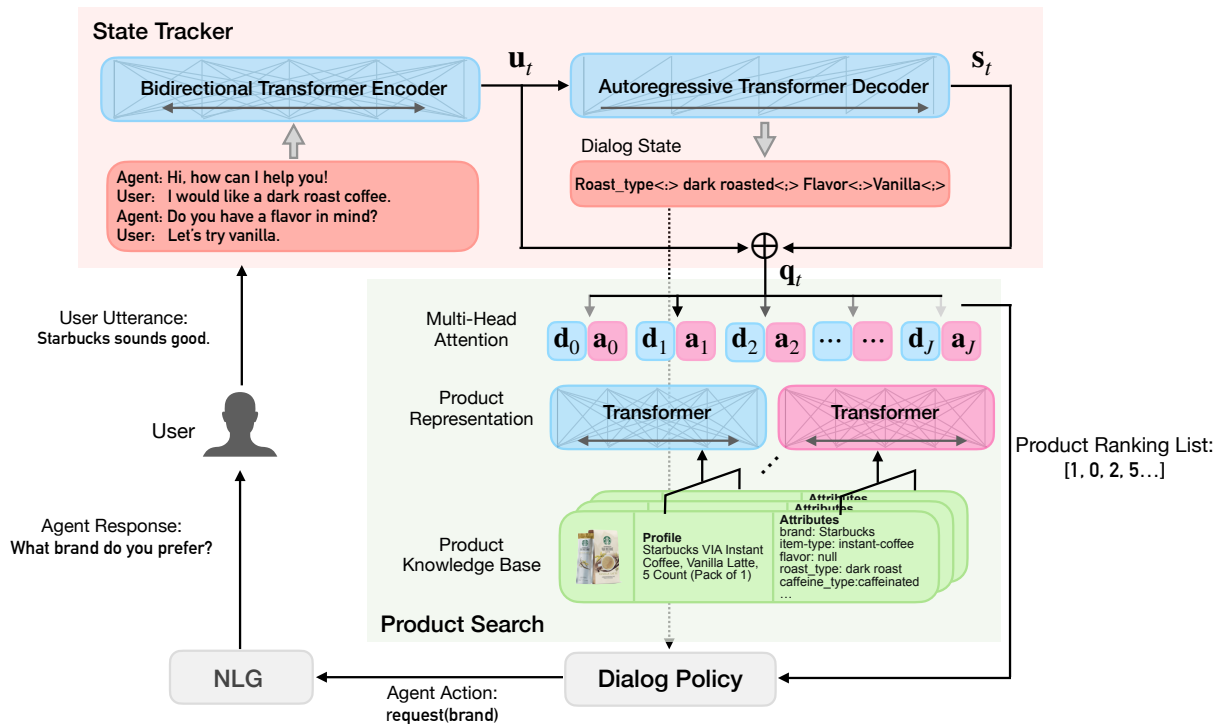


Figure 1: Illustration of our end-to-end conversational search system. State Tracker module takes utterances to predict dialog state S_t using the seq-to-seq transformer. Product Search module matches products represented by transformers with query representations q_t using a multi-head attention mechanism. Dialog Policy module takes inputs from S_t , intent and ranked product list, and decides the responses. NLG module composes system responses as instructed by Dialog Policy, and displays them to user.

it hard for utterance understanding model training. The approach adopted by Yan et al. (2017) mines phrases of shopping from community sites and uses crowd-sourcing to label utterance intents. Although costing less labors, this work did not construct the full dialogs. As a comparison, in this paper we constructed the full shopping search dialogs through real user behavior data, with user utterances filled by transferring from existing dialogs of similar domains.

3 End-to-End Conversational Search with Multi-task Learning

Our conversational search system, depicted in Figure 1, consists of four major modules: State Tracker (ST), Product Search, Dialog Policy, and Natural Language Generation (NLG). The State Tracker module interprets the dialog content, outputs user intent, along with product attributes that the user is interested in. The Product Search module returns a list of products which are matching with tracked attributes user interested in. Based on output from State Tracker, the Dialog Policy manages agent response according to user intents and the candidate

search result, and the NLG module transforms the response into natural language displayed to user.

To address the challenge of imperfect product attribute schema/knowledge, our Product Search module leverages both structured product attributes and unstructured text. To mutually benefit from each other’s learning, we integrate State Tracker and Product Search together through multi-task learning, and build an end-to-end trainable search system.

3.1 State Tracker

Unlike previous work that treats state tracking as a multi-label classification task (Zhu et al., 2020; Wen et al., 2017a), we redefine the state tracking task as a sequence-to-sequence problem. As shown in Figure 1, we link the slots and values of dialog state with special delimiter tokens, turning it into a sequence. Then we employ a transformer network to translate dialog turns into state, which encodes the dialog lines with a bidirectional transformer encoder and generate state sequence autoregressively.

At each turn in the dialog, the State Tracker module outputs 1) the dialog *state* S and 2) the user utterance *intent* I , where S are attribute-values

grouped by product attributes, representing the system’s tracking of user’s preferred search criteria, and *intent* $I \in \mathbb{I}$ is an enumerable value from $\mathbb{I} = \{request, inform, ask_attribute_in_n, buy_n\}$.

Formally, given a dialog at turn t , we have all history records of $\{R_0, U_0, R_1, U_1, \dots, R_t, U_t\}$, where R_t and U_t are *system response* and *user utterance* at turn t respectively. We then use a transformer model (Devlin et al., 2019) to predict the string S_t , the *state* at t :

$$S_t = trans(concat(R_0; U_0; \dots; R_{t-1}; U_{t-1})) \quad (1)$$

where $trans(\cdot)$ is the transformer, and $concat(\cdot)$ is a string concatenation function. For *state* prediction, we use the loss function:

$$\mathcal{L}_s = - \sum_t \sum_i \log P(y_t^{i*}) \quad (2)$$

where y_t^{i*} denotes the ground-truth value for i -th item of output sequence at t -th turn.

We also use an MLP layer to predict the *intent* at t :

$$\mathbf{z}_t = \mathbf{W}_I \cdot \text{ReLU}(\mathbf{u}_t) + b_I \quad (3)$$

$$P_t(I_i) = \text{softmax}(\mathbf{z}_t) \quad (4)$$

where \mathbf{u}_t is the mean pooling of last layer output from the encoder in Equation (1), \mathbf{W}_I and b_I are trainable parameters, and $P_t(I_i)$ represents the likelihood of *intent* $I_i \in \mathbb{I}$ from user utterance at t -th turn. We use the following loss function for intent prediction:

$$\mathcal{L}_I = - \sum_t \log P_t(I_t^*). \quad (5)$$

where I_t^* is the ground truth of *intent* at turn t .

3.2 Product Search

At each turn t , given current *state*, Product Search module estimates the matching likelihood $P_t(p_j)$ for each product p_j , and then rank the products to be displayed to user (Figure 1).

Query Representation We represent the product query as $\mathbf{q}_t = \mathbf{u}_t \oplus \mathbf{s}_t$, where \mathbf{s}_t is the state representation obtained by mean pooling the last layer of decoder in Equation (1), and \oplus denotes the vector concatenation operator.

Product Embedding We represent j -th product with $\mathbf{p}_j = \mathbf{d}_j \oplus \mathbf{a}_j$, where \mathbf{d}_j is the mean pooling of the last encoding layer of: $trans(\text{description text of } p_j)$, and \mathbf{a}_j is the product attribute embedding. In particular, we obtain \mathbf{a}_j by mean pooling the last layer of: $trans(\text{attribute sequence of } p_j)$, where the attribute sequence is constructed as state sequence.

The introduction of profile embedding \mathbf{d}_j compensates for the missing matching clue when product schema is incomplete or attribute values are missing, since they may be extracted from product text.

Search with Multi-Head Attention We use multi-head attention mechanism to match query and products. At dialog turn t , we first calculate a product context vector \mathbf{head}_t^k based on the glimpse operation (Vinyals et al., 2016):

$$\mathbf{a}_t^j = \mathbf{v}_s^{k\top} \tanh(\mathbf{W}_p^k \mathbf{p}_j + \mathbf{W}_q^k \mathbf{q}_t) \quad (6)$$

$$\alpha_t = \text{softmax}(\mathbf{a}_t) \quad (7)$$

$$\mathbf{head}_t^k = \sum_t \alpha_t^j \mathbf{W}_p^k \mathbf{p}_j \quad (8)$$

where α_t are attention weights, and $\mathbf{W}_p^k, \mathbf{W}_q^k, \mathbf{v}_s^k$ are trainable parameters for head k . We then concatenate K attention heads each with individual parameter sets, $\mathbf{head}_t = \bigoplus_{0 \leq k \leq K} \mathbf{head}_t^k$.

We then form likelihood of product p_j at t -th turn as:

$$\mathbf{e}_t^j = \mathbf{v}_p^\top \tanh(\mathbf{W}_p' \mathbf{p}_j + \mathbf{W}_h \mathbf{head}_t) \quad (9)$$

$$P_t(p_j) = \text{softmax}(\mathbf{e}_t) \quad (10)$$

where $\mathbf{v}_p, \mathbf{W}_p'$, and \mathbf{W}_h are trainable parameters. We use the following loss function for product search job:

$$\mathcal{L}_p = - \sum_t \log P_t(p_t^*) \quad (11)$$

where p_t^* is the ground-truth of product. Finally, we rank products with their likelihood, and return top products to Dialog Policy module for displaying.

3.3 Multi-task Learning

Our end-to-end training links all three tasks (*state* prediction, *intent* prediction and Product Search) together through multi-task learning:

$$\mathcal{L} = \alpha \mathcal{L}_s + \beta \mathcal{L}_I + \gamma \mathcal{L}_p \quad (12)$$

where α, β, γ are tunable hyper-parameters. With multi-task learning, these three tasks can enhance each other with shared weights and back-propagated errors.

The training data requires intent and attribute annotation for each utterance, and purchased products with product attributes and text profiles (optional) associated with each dialog.

3.4 Dialog Policy and Natural Language Generation

During the conversation, the agent needs to propose additional attributes for user to narrow down the search. When triggered, we filter our product knowledge base using current *state S* to retrieve products matching with the criteria, then use EMDM (Entropy Minimization Dialog Management) (Wu et al., 2015) to select the proposed attribute with maximum entropy among filtered products, and show user recommended narrowing down question.

The Natural Language Generation module translates the action decision from the Dialog Policy module to natural language, e.g. *request(brand)* \rightarrow *Do you have a brand in mind?*. In this paper we simply use manually written agent templates.

4 Dialogue Dataset Construction

We address the challenge of lack of conversational shopping search training data by proposing **M2M-UT**, a method that automatically constructs dialog datasets. Unlike previous works (Saha et al., 2018) that rely on crowd-source to generate utterances, M2M-UT can automatically generate utterances with transfer.

We hypothesize that the conversation between the user and the shopping agent is guided by customer’s *intents* that 1) span user’s utterance in natural language, and 2) change according to agent’s responses. Therefore, our dataset construction has two steps: 1) we use utterance transfer (UT) to generate utterances from existing dialog datasets of similar domains, and 2) we generate the outline of dialog using customer browsing records using Machine Talking To Machine (M2M) Saha et al. (2018).

4.1 Utterance Generation by Transfer

For utterance generation, widely used methods such as WoZ and M2M still require workers to create the various utterances, and thus are not easy

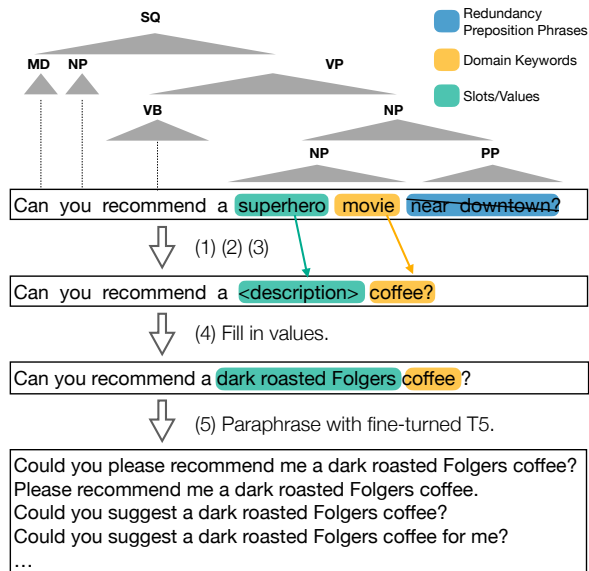


Figure 2: Utterance generation algorithm to generate variant utterances for coffee shopping domain. The utterance example employed in this figure is from MDC dataset (Li et al., 2018b). An utterance is first transferred to our domain with the help of constituency parser and then paraphrased to enhance the variance.

to scale up in the shopping conversation application. We found that dialogues from existing task-oriented domains such as movie ticketing or restaurants reservation contain rich form of utterances similar to shopping, for example, “... sounds good” is seen from both movie ticketing and shopping conversations. We propose utterance transfer (UT), a novel approach that generates shopping utterances from related task-oriented domains.

As shown in Figure 2, UT consists of five stages. (1) **remove redundant phrases**: we remove the redundant phrases that not commonly seen in online shopping (e.g. location and time) with syntax rules. We employ a constituency parser (Kitaev and Klein, 2018) to get the syntax tree of the sentence and remove the PPs (preposition phrases) and NPs (noun phrases) referring to location and time. (2) **replace values with slots**: we identify and replace values with slots according to the original dataset annotations. For example, in Figure 2, we identify value “superhero” using the annotation, and replace it as slot “<description>”. This step turns a complete utterance into a template. (3) **keyword replacement**: we replace verbs and nouns with those from online shopping domain with rules, e.g. “movie” to “coffee” and “watch” to “drink”. (4) **fill slots**: we fill the slots with values according to

Search Behavior Data	
Search keywords: vanilla instant coffee packets Goal: {flavor:vanilla, item_type:instant-coffee, brand:Folgers, roast_type=medium roast,...} Purchase: ID=B074FLFKNV	
Outline	Utterances
S: greeting()	Hello, what can I do for you?
U: inform(flavor=vanilla, item_type=instant-coffee)	Please find me vanilla instant coffee packets.
S: request(brand)	Do you have a brand in mind?
U: inform(brand=Folgers)	Let's try Folgers.
S: push(top_5)	I found you following products: <Products List>
U: ask_attr_in_n(roast_type, index=2)	What roast type is it in the second image.
S: inform(roast_type=medium roast)	It is medium roast.
U: buy_n(index=2)	I will buy the second one.
S: notify_success()	Your order has been placed.

Table 1: M2M-UT dialog generation. M2M-UT first generates the dialog outline with search behavior data, then translates it to utterances with the method illustrated in Figure 2. S and U denote System and User respectively. We use search keywords to generate the first user utterance.

user’s action. (5) **paraphrase**: to augment the diversity of utterance, we use a fine-tuned T5 model (Raffel et al., 2020) to paraphrase the utterance.

Paraphrase One pitfall of utterances generated by templates and rules are the lack of diversity, whereas real conversations usually contain various ways of expressing the same intents. As paraphrase can improve the performance of dialog system (Gao et al., 2020), we employ a pre-trained neural paraphrase model to augment the variance of templates. Specifically, we use a T5 model (Text-to-Text Transfer Transformer)⁵ (Raffel et al., 2020) that is fine-tuned on paraphrase dataset, Quora Question Pairs⁶.

4.2 Dialog Generation

Our online shopping dialog in conversational search is supported by the dialog *outlines*, which consists of *intent* and its parameters. For user utterance intents as shown in Table 1, their parameters are typically a list of product attributes with their values. For agent intents, parameters are either attribute values, or operation parameters that agent should execute with (e.g., push(top_5)). Similar to the dialog system presented in Section 3, we use

⁵<https://github.com/ramsrigouthamg/Paraphrase-any-question-with-T5-Text-To-Text-Transfer-Transformer>

⁶<https://www.quora.com/q/quoradata/First-Quora-Dataset-Release-Question-Pairs>

state to track agent’s understanding of user’s search criteria.

We use real e-commerce search behavior data to supervise the construction of intent flow in the dialog. Each anonymous search session contains a query and the final purchased product. We first extract product attribute values from the search keywords as the initial attribute customer interested in, i.e., *initial state*. We then follow M2M (Machine Talking to Machine) (Shah et al., 2018) to generate the transition of dialogue outlines turn by turn. M2M runs in a self-playing manner by simulating the dialog with a user simulator and a system agent. We build an agenda-based user simulator initialized by the search behavior data, and use a finite state machine (Hopcroft et al., 2007) as the system agent.

By comparing *initial state* and the finally purchased product, we find that users were not always aware of the full search criteria at the beginning, therefore the dialog is constructed to simulate how agent helps user to fill the gap through attribute refinement. Specifically, as shown in Table 1, user starts with *initial state* (e.g., flavor=vanilla). Given current *state*, agent in the next turn proposes a new attribute (e.g., brand) using the policy EMDM (Entropy Minimization Dialog Management) (Wu et al., 2015) to narrow down the search. User’s response in the next turn will be based on attribute value of the purchased product (e.g., brand=Folgers), which also updates the *state*. Then agent displays a list of products in the next turn (e.g., push(top_5)). If purchased product appear in push list, user asks questions, commits purchase, and ends the dialog (successful). Otherwise agent proposes a new attribute, and continues the conversation. Dialog ends when length exceeds 20 (unsuccessful).

We finally translate the generated outline into natural language by using corresponding utterance templates generated after step (3) in Section 4.1, and finalize the utterance following step (4) and (5) in Section 4.1. After these steps, we have a complete shopping search dialog.

5 Experiments

5.1 Datasets

Our dataset includes three parts: user search behavior data, dialogs, and product knowledge base.

The user search behavior data is a collection of user search keywords and their final purchased

Metric	CSD-UT
#Dialogs	49,999
#Total utterances	942,766
#User utterances	471,383
Avg. #Turns per dialog	18.85
Avg. #Tokens per utterance	6.57

Table 2: Statistics of generated dialog dataset.

Metric	Product KB.
#Products	154,161
#Attributes	13
Avg. #Values per attribute	1111.13
Avg. #Tokens per profile	17.34
Vacant ratio of values	32.16%

Table 3: Statistics of product knowledge base.

products sampled from e-commerce platform. We applied the dialog generation method described in Section 4 on the coffee shopping domain. We leveraged the utterances from dataset MDC (Li et al., 2018b) and MMD (Saha et al., 2018) and transferred 4 intents from their domains (i.e. movie ticketing, restaurant reservation, fashion shopping), which generated 49,999 dialogs, with each of the dialog contains on average 18.85 turns (Table 2). In addition, we built up a gold-standard test set of 196 dialogs manually written by workers to evaluate the performance.

For the product knowledge base, we sampled 154,161 coffee products from the e-commerce platform. As shown in Table 3, each product has a text profile with average 17.34 tokens and also the attribute-value pairs for 13 different attributes. The vacant ratio of values is 32.16%, which indicates potential missing attribute values for products.

5.2 Settings

Hyper-parameters All the transformers used in experiment have 4 sublayers with hidden size of 256, and a word2vec (Mikolov et al., 2013) of 256 dimension is trained to initialize the embedding matrix. Our model used a vocabulary of 50257 entries for text embedding, and 14700 entries for attribute embedding. The models in experiments were trained with AdamW (Loshchilov and Hutter, 2017) optimizer with the initial learning rate of $1e-4$ and batch size of 16. The initial learning rate is selected based on validation loss. We used learning rate scheduler to cut the learning rate by half every time the performance drops and stop training once the performance has three straight drops. Our

Model	SR@5	SR@10
TC-bot (Li et al., 2017)	35.71	51.02
ConvLab-2 (Zhu et al., 2020)	44.89	54.08
ConvSearch		
w/ Rule Search	39.79	50.51
w/ Neural Search (attr.)	46.42	57.14
w/ Neural Search (text.)	48.47	59.69
w/ Neural Search (attr. & text.)	51.53	61.22

Table 4: Evaluation of the end-to-end system. attr. and text. denote attribute and product text respectively. The best score per metric is in bold. Our model outperforms the competitors by 6.64%. Rule search employs direct attribute matching as traditional work.

model was trained on a Nvidia Tesla P100 machine with 16G memory, and the strongest model (ConvSearch w/ Neural Search (attr.&text.)) took 35 hours for convergence. For multi-task learning, we briefly set α, β, γ as 1. To save memory, we let the encoder of state tracker and encoder of profile share the parameters, and employed tf-idf to narrow down the search space into 400 products for product search module.

Evaluation Metrics We use the success rate ($SR@t$) to measure the ratio of successful conversations, i.e. recommended the ground-truth item in t turns. We set the max turn t of a session to 5 or 10 and standardized the recommended list length as 5. Besides, we used recall, precision and F1 to evaluate the performance of state prediction, and reported the performance on slot and value respectively.

Baselines For state tracking task, we compared against the following baselines: e2e-Trainable (Wen et al., 2017b) which encodes utterances with a convolutional neural network (CNN), ZS-DST (Rastogi et al., 2020), a Bert-based model which first judges the presence of each slot then the start and end location. We also constructed a baseline by replacing transformers in our system with one-layer LSTMs. For the end-to-end system, we compared against two baselines: TC-bot (Li et al., 2017), a modularized neural dialogue systems for task-completion, and ConvLab-2⁷ (Zhu et al., 2020), an open-source toolkit for building, evaluating, and diagnosing a task-oriented dialogue system.

5.3 End-to-End System Evaluation

Table 4 shows the end-to-end task (success rate) comparisons, where our method outperforms base-

⁷we employ the strongest setting they reported, BERTNLU+RuleDST+RulePolicy+TemplateNLG.

Model	State-Attr			State-Value			Intent		
	R	P	F1	R	P	F1	R	P	F1
e2e-Trainable (Wen et al., 2017b)	92.67	82.74	87.46	90.98	86.57	88.66	95.75	95.91	95.82
ZS-DST (Rastogi et al., 2020)	96.97	89.55	93.11	91.41	87.70	89.51	96.43	97.89	97.15
LSTM + classification	92.34	89.72	91.01	88.97	82.31	85.51	95.65	94.26	94.94
State Tracker w/o Search	97.53	93.29	95.36	92.01	87.27	89.58	99.73	99.68	99.70
State Tracker w/ Search	98.15	93.41	95.72	93.15	87.44	90.20	99.70	99.69	99.69

Table 5: Evaluation of state tracking task. R and P denote recall and precision.

Model	R	P	F1
tf-idf	5.58	1.16	1.86
Product Search w/attr.	15.53	3.10	5.16
Product Search w/text.	19.27	4.84	7.74
Product Search w/text. & attr.	26.20	5.47	9.05

Table 6: Independent evaluation of search task. This experiment shows the benefit of combining product text profile and attribute for search. attr. is abbreviation for product attribute. The best score per metric is in bold.

lines significantly by 6.64%. This indicates the effectiveness of our end-to-end framework that deeply combines the dialog and search system, while ablation studies (last three rows in Table 4) also justify that leveraging both product text and attribute performs better than using only one of them.

5.4 State Tracker Evaluation

Table 5 shows the performance comparisons of state tracking task. It shows that our method outperforms all baselines in both *state* prediction and *intent* prediction tasks, which is because our state tracking task can better embed the context by concatenating the language of turns together. We also found State Tracker alone without Product Search task showed lower performance, suggesting the effectiveness of multi-task learning.

5.5 Product Search Evaluation

Table 6 shows ablation studies of the Product Search module, along with comparisons with a simple tf-idf baseline. In particular, after the 3rd turn of dialog, we selected top-5 products with highest probability from the list returned by Product Search module, and calculated recall, precision and F1 value against the ground-truth purchased product. We can see that the end-to-end search significantly improved the search recall by 4.69 times over the tf-idf baseline. Improvement induced by combining text and attribute embedding suggests the benefits of combining product text and attributes in search task.

Dataset	Attr	Value	Intent
UT w/o Syntax & Paraphrase	72.53	65.98	88.53
UT w/o Syntax	85.71	79.61	98.14
UT w/o Paraphrase	87.51	75.12	97.55
UT	95.72	90.20	99.69

Table 7: The effectiveness of utterance generation methods for utterance understanding. The numbers in the table are F1 scores. We can see that both syntax and paraphrase improve the dialog data quality.

Method	Coherence	Fluency	Approp.
TC-bot	2.98	3.42	3.20
ConvSearch	3.58	3.54	3.66

Table 8: Human Evaluation Result. Approp. is short for Appropriateness.

5.6 Dialog Generation Method Evaluation

We next conducted ablation studies on the data construction method. We evaluated the effectiveness of each component using the performance of State Tracker task. For each configuration in Table 7, we trained the State Tracker module with corresponding dataset, and reported the performance on a manually prepared test set. As shown in the table, the module performance degrades without syntax analysis since redundant phrase (e.g. time, location) won't be removed from the utterance. Similarly, module performance degrades without paraphrase since language variance will be weakened. These suggest that both removing redundancy with syntax and increasing variance with paraphrase are effective to improve the training dataset quality.

5.7 Human Evaluation

We also performed human evaluations on system responses. For each method, we collected 100 dialogs and asked three workers to evaluate them with three metrics: coherence, fluency and appropriateness. All metrics have five grades: from 1(worst) to 5(best), where 3 denotes 'good'. As shown in Table 8, ConvSearch outperforms the baseline model in all three metrics.

6 Conclusion and Future Work

In this work, we built an end-to-end conversation search system for online shopping, where we deeply combined the dialog and search system with multi-task learning. In particular, our product search module leverages both product attribute and text to retrieve products, which mitigates the imperfect product schema/knowledge challenges. To address issue of lacking in-domain dialog dataset, we proposed a dataset transfer method and constructed shopping dialog dataset from user search behavior data and existing dialogs of similar domain. The proposed dataset construction method lowers the cost, making it possible to scale-up to broader use scenarios.

We will leave it to future work to expand the methodology across more shopping categories, and broader use scenarios such as clinical conversations and customer service, etc.

Acknowledgments

Hao He is supported by the National Key Research and Development Program of China under Grant 2018YFC0830400, the Basic Research Project of Shanghai Science and Technology Commission under ECNU-SJTU joint Grant 19JC1410102, the Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102. This research is partially supported by the Shanghai Science and Technology Innovation Action Plan under Grant 20511102600.

References

- Keping Bi, Qingyao Ai, Yongfeng Zhang, and W. Bruce Croft. 2019. [Conversational product search based on negative feedback](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 359–368. ACM.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Silin Gao, Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. [Paraphrase augmented task-oriented dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 639–649. Association for Computational Linguistics.
- John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. 2007. *Introduction to automata theory, languages, and computation, 3rd Edition*. Pearson international edition. Addison-Wesley.
- J. F. Kelley. 1984. [An iterative design methodology for user-friendly natural language office information applications](#). *ACM Trans. Inf. Syst.*, 2(1):26–41.
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Melbourne, Australia. Association for Computational Linguistics.
- Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. [Estimation-action-reflection: Towards deep interaction between conversational and recommender systems](#). In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 304–312. ACM.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018a. [Towards deep conversational recommendations](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 9748–9758.
- Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Çelikyilmaz. 2017. [End-to-end task-completion neural dialogue systems](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017 - Volume 1: Long Papers*, pages 733–743. Asian Federation of Natural Language Processing.
- Xiujun Li, Panda Sarah, Jingjing Liu, and Jianfeng Gao. 2018b. [Microsoft dialogue challenge: Building end-to-end task-completion dialogue systems](#). In *SLT*.
- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *CoRR*, abs/1711.05101.
- Kai Luo, Scott Sanner, Ga Wu, Hanze Li, and Hojin Yang. 2020. [Latent linear critiquing for conversational recommender systems](#). In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 2535–2541. ACM / IW3C2.
- Julian J. McAuley, Rahul Pandey, and Jure Leskovec. 2015. [Inferring networks of substitutable and complementary products](#). In *Proceedings of the 21th*

- ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, pages 785–794. ACM.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- 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](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8689–8696. AAAI Press.
- Amrita Saha, Mitesh M. Khapra, and Karthik Sankaranarayanan. 2018. [Towards building large scale multimodal domain-aware conversation systems](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 696–704. AAAI Press.
- Pararth Shah, Dilek Hakkani-Tür, Bing Liu, and Gökhan Tür. 2018. [Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 3 (Industry Papers)*, pages 41–51. Association for Computational Linguistics.
- Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. 2016. [Order matters: Sequence to sequence for sets](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve J. Young. 2017a. [Latent intention dialogue models](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3732–3741. PMLR.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina Maria Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve J. Young. 2017b. [A network-based end-to-end trainable task-oriented dialogue system](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, pages 438–449. Association for Computational Linguistics.
- Ji Wu, Miao Li, and Chin-Hui Lee. 2015. [A probabilistic framework for representing dialog systems and entropy-based dialog management through dynamic stochastic state evolution](#). *IEEE ACM Trans. Audio Speech Lang. Process.*, 23(11):2026–2035.
- Zhao Yan, Nan Duan, Peng Chen, Ming Zhou, Jianshe Zhou, and Zhoujun Li. 2017. [Building task-oriented dialogue systems for online shopping](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4618–4626. AAAI Press.
- Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. [Towards conversational search and recommendation: System ask, user respond](#). In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, October 22-26, 2018*, pages 177–186. ACM.
- Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takanobu, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020. [Convlab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2020, Online, July 5-10, 2020*, pages 142–149. Association for Computational Linguistics.