

A Speaker Turn-Aware Multi-Task Adversarial Network for Joint User Satisfaction Estimation and Sentiment Analysis

Kaisong Song^{1,2}, Yangyang Kang^{1*}, Jiawei Liu^{3*}, Xurui Li¹, Changlong Sun¹, Xiaozhong Liu⁴

¹ Alibaba Group, Hangzhou, China

² Northeastern University, Shenyang, China

³ Wuhan University, Wuhan, China

⁴ Worcester Polytechnic Institute, Worcester, USA

kaisong.sks@alibaba-inc.com, yangyang.kangyy@alibaba-inc.com, laujames2017@whu.edu.cn

xurui.lxr@alibaba-inc.com, changlong.scl@taobao.com, xliu14@wpi.edu

Abstract

User Satisfaction Estimation is an important task and increasingly being applied in goal-oriented dialogue systems to estimate whether the user is satisfied with the service. It is observed that whether the user’s needs are met often triggers various sentiments, which can be pertinent to the successful estimation of user satisfaction, and vice versa. Thus, User Satisfaction Estimation (USE) and Sentiment Analysis (SA) should be treated as a joint, collaborative effort, considering the strong connections between the sentiment states of speakers and the user satisfaction. Existing joint learning frameworks mainly unify the two highly pertinent tasks over cascade or shared-bottom implementations, however they fail to distinguish task-specific and common features, which will produce sub-optimal utterance representations for downstream tasks. In this paper, we propose a novel Speaker Turn-Aware Multi-Task Adversarial Network (STMAN) for dialogue-level USE and utterance-level SA. Specifically, we first introduce a multi-task adversarial strategy which trains a task discriminator to make utterance representation more task-specific, and then utilize a speaker-turn aware multi-task interaction strategy to extract the common features which are complementary to each task. Extensive experiments conducted on two real-world service dialogue datasets show that our model outperforms several state-of-the-art methods.

Introduction

The proliferation of complaint and impolite language within service dialogues can create salient social and financial issues, such as marketing harassment and brand reputation damage. Many platforms have built their serving rating systems that rate the suitability and quality of online service. For example, over 77% of buyers on Taobao¹ communicated with sellers before placing an order (Gao and Zhang 2011). User satisfaction estimation (USE) is an important yet challenging task, which has drawn great attention from both industries and research communities (Song et al. 2019; Liu et al. 2021b; Higashinaka et al. 2019; Deng et al. 2022).

*Corresponding authors

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.taobao.com/>

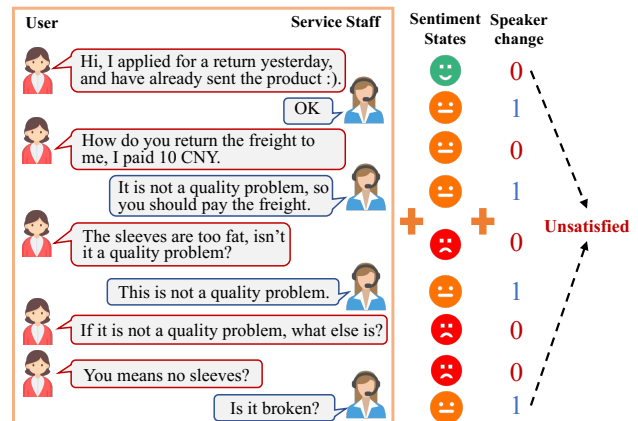


Figure 1: A typical serving rating system and an unsatisfied service dialogue with sentiment and speaker turn change labels. Graphic emoji denotes positive, neutral or negative. Label 1 or 0 denotes whether speaker turn changes or not.

Intuitively, whether the user is satisfied with a service staff largely depends on the fulfillment of the user’s needs. The workflow of a typical service rating system is displayed in Figure 1. The user interacts with the online service staff in the free-form multi-turn dialogue, which is the basic service mode on E-commerce platform. A complete dialogue will be input into the rating system for satisfaction rating. The user expresses positive sentiments in the beginning and interacts with the service staff by using a textual smiley face “:).” As the conversation progresses, sentiment states of the user gradually change from positive to negative, and the user is ultimately very unsatisfied and fails to reach an agreement with the service staff on the return issue. Automatically detecting such low-quality and unsatisfying service is important. For the retail shopkeepers, they can quickly locate such service dialogue and find out the reason to take remedial actions, such as “user return visit” and “staff service improvement”. For the platforms, by detecting and analyzing such cases, they can define clear-cut rules, say “not fitting well is not a quality issue, and the buyers should pay the freight.”

User satisfaction often relates to many linguistic and sta-

tistical properties from content and syntax (Shan et al. 2019; Song et al. 2020; Liu et al. 2021a), speaker profiles (Yao et al. 2020; Abraham and Choe 2021) to sentiment (Gadri et al. 2021; Song et al. 2021), act (Deng et al. 2022) and deep semantics (Sun et al. 2021; Kachuee et al. 2021), which are regularly considered as key challenges in semantic understanding of dialogues. Different from previous studies that focus on modeling various properties, we find the sentiment states of the speakers provide rich clues for explicitly assisting with user satisfaction estimation. Without guidance of utterance-level sentiment clues, user satisfaction estimation may become very difficult to understand semantically for long dialogues. Multi-task learning (MTL) allows two or more tasks to be learned jointly, thus sharing information and features between the tasks (Caruana 1997). In this work, we focus on the dialogue-level user satisfaction estimation; hence we refer to it as the *main task*, while the task that is used to provide additional sentiment knowledge, i.e., utterance-level sentiment analysis, is referred to as the *auxiliary task*. Thus, we aim to capture the sentiment feature resulting from a joint learning setup through shared parameters. This will encompass the sentimental content of utterances that is likely to be predictive of user satisfaction.

Earlier MTL studies on USE often resort to implicit or explicit sentiments, which can not resolve correlated interference between task-private feature and shared feature well. Song et al. (2019) aggregated the predictive utterance-level latent sentiments as the estimation of dialogue-level user satisfaction via multiple instance learning (Angelidis and Lapata 2018). However, such cascade implementation will drag performance greatly because sentiment features and satisfaction features will interfere with each other. Liu et al. (2021b) considered users' negative sentiments as unsatisfied signals to handover chatting from machine to human. Their shared-bottom implementation alleviates previous problem by simply introducing separate shallow fully-connected layers to derive task-specific utterance features from the same inputs. Despite the effectiveness of incorporating sentiment clues on USE in previous works, there are several issues that remain to be tackled. (i) The performance of task models is influenced by the task-specific input features. Because the shared utterance features can exist in the task private space and the task-specific features creep into the shared space (Zhou et al. 2021). (ii) The task models should interact with each other via soft-parameter sharing instead of hard-parameter sharing. Because the later will produce sub-optimal utterance representation for downstream tasks. (iii) Turn-taking modeling across different speakers is necessary, which determines whether sentiment and topics of the consecutive utterances are consistent. For instance, given a *negative* utterance from a user, if the following utterance is from a service staff, then the following sentiment is likely to be *positive* or *neutral*; however, if there is no change in speakers, then the sentiment is probably still *negative*.

In this paper, we address the joint user satisfaction estimation and sentiment analysis by proposing a novel and extensible Speaker Turn-Aware Multi-Task Adversarial Network (STMAN) model, which contributes in three ways:

- We propose a novel multi-task adversarial strategy which

uses a task discriminator to differentiate task-specific utterance features explicitly by classifying each utterance into different task categories, i.e, USE and SA.

- We propose a novel speaker-turn aware interaction strategy which can extract the common sentiment features of each utterance while each task can still learn its task-specific semantic and speaker-turn change features.
- Extensive experiment conducted on two real-world dialogue datasets validate the effectiveness of our approach.

Related Work

User Satisfaction Estimation. User satisfaction, which is related to fulfillment of a specified desire or goal, is essential in evaluating and improving user-centered goal-oriented dialogue systems. Song et al. (2019) studied the question-and-answer utterance matching within dialogues and proposed a context assisted multiple instance learning model to predict user satisfaction. Yao et al. (2020) conducted a more in-depth exploration of the interactive relationship of multiple rounds of human-computer interaction. They introduced personal historical dialogue to model satisfaction bias. However, both the methods predict satisfaction solely based on content features. Therefore, several efforts have been made on modeling user satisfaction from temporal user behaviors or actions when interacting with the systems (Mehrotra et al. 2019; Liu et al. 2021b; Deng et al. 2022). Mehrotra et al. (2019) extracted informative and interpretable action sequences (e.g., Click, Pause, Scroll) from user interaction data to predict user satisfaction towards the search systems. Liu et al. (2021b) conducted machine-human chatting handoff and service satisfaction analysis jointly. Deng et al. (2022) leveraged the sequential dynamics of dialogue acts to facilitate USE in goal-oriented conversational systems via a unified joint learning framework. However, these methods utilize task-specific features in a rough way, thus producing sub-optimal utterance representation for downstream tasks.

Dialogue Sentiment Analysis. This task assigns a proper sentiment label to each utterance within any dialogue. Wang et al. (2020) learned topic enriched utterance representations for utterance-level sentiment classification. Qin et al. (2021) treated an utterance as a vertex and added an edge between utterances of the same speakers to construct cross-utterances connections; such connections are based on specific speaker roles. Majumder et al. (2019) kept track of the individual party states throughout the conversation and used this information for emotion classification. However, such personalized models have low generalizability because they are severely limited by large amounts of learnable person-related parameters. Previous methods incorporate speaker information by proposing more complex and specialized models, which inevitably introduce a large number of parameters to train. Recently, He et al. (2021) considered speaker-aware turn changes by concatenating turn embedding and utterance embedding together in a rough way. Here, we utilize sentiment states of speakers to assist with user satisfaction estimation, which is rarely studied before.

Multi-Task Learning. USE and SA can be studied in a joint framework. Bodigutla et al. (2020) proposed to jointly

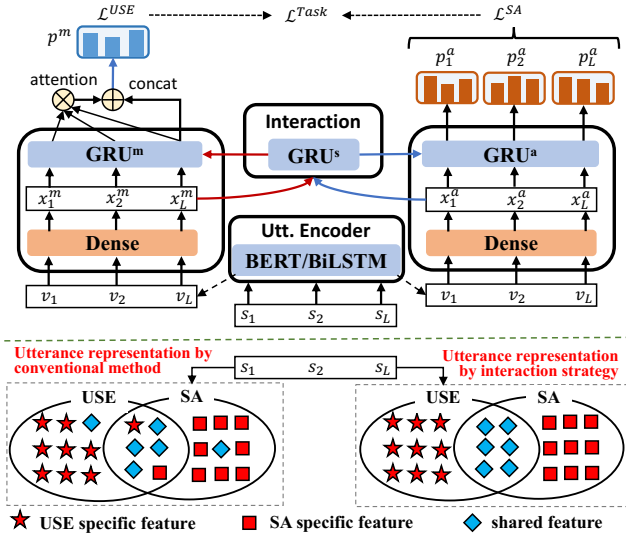


Figure 2: The architecture of Basic model.

predict turn-level labels and dialogue-level ratings, which used a BiLSTM variant to weight each turn’s contribution towards the ratings. Ma, Gao, and Wong (2018) proposed a joint framework that unifies two highly pertinent tasks. Both tasks are trained jointly using weight sharing to extract the common and task-invariant features. However, these frameworks can not be directly adapted to our dialogue scenario well. Firstly, the same dialogue inputs should be treated in a more task-specific way. Secondly, speaker turn change information is ignored, which determines whether sentiment and topics of the consecutive utterances are consistent. Finally, features should be shared in a more controllable way.

Basic Multi-Task Learning Network

Figure 2 shows the overall architecture of our basic model, which consists of four components. This paves way for introducing our fully configured model.

Shared Utterance Encoder

Let any dialogue \mathcal{D} as $[s_1, s_2, \dots, s_L]$, where L is the number of consisted utterances. The shared utterance encoder Enc is established to map each utterance $s_t \in \mathcal{D}$ to a latent representation: $\mathbf{v}_t = Enc(s_t)$, where function Enc can be a standard BiLSTM (Schuster and K.Paliwal 1997) with attention mechanism (Song et al. 2019; Liu et al. 2021b) or pre-trained language models, like BERT (Devlin et al. 2019). Thus, the dialogue \mathcal{D} can be represented as a list of utterance vectors, i.e., $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_L]$. Here, we omit the details of the shared utterance encoder due to limited space.

Multi-Task Interactive Layer

The performance of task models is influenced by the task-specific features. We first apply two dense layers over the shared utterance vectors \mathbf{V} respectively to make them task-specific, which can be formulated as below:

$$\mathbf{X}^m = Dense^m(\mathbf{V}) \quad , \quad \mathbf{X}^a = Dense^a(\mathbf{V}) \quad (1)$$

where vectors $\mathbf{X}^m = \{\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_L^m\}$ and $\mathbf{X}^a = \{\mathbf{x}_1^a, \mathbf{x}_2^a, \dots, \mathbf{x}_L^a\}$, superscripts “ m ” and “ a ” denote the *main* (USE) and *auxiliary* (SA) tasks, respectively. To model the correlations between the separate tasks, we adopt a multi-task interaction strategy which consists of two hidden layers for each task: one is used to extract the common patterns via the shared parameters, and the other is used to capture task-specific features via the separate parameter sets. Accordingly, each task is assigned a shared GRU layer and a task-specific GRU layer, which hopefully can be used to capture the shared and task-private representations for different tasks. For the *main task*, the USE-specific features \mathbf{X}^m are fed into a soft-parameter sharing standard GRU^s unit (Cho et al. 2014). Meanwhile, we use a variant GRU^m unit for producing USE-specific utterance representations based on \mathbf{X}^m and outputs of GRU^s . The GRU^s and GRU^m units are formulated as below:

$$\begin{aligned} \mathbf{h}_t^s &= GRU^s(\mathbf{h}_{t-1}^s, \mathbf{x}_t^m) \\ \mathbf{h}_t^m &= GRU^m(\mathbf{h}_{t-1}^m, \mathbf{h}_t^s, \mathbf{x}_t^m) \end{aligned} \quad (2)$$

where the hidden output $\mathbf{h}_t^m \in \mathbb{R}^K$ of the GRU^m unit at step t is dependent on the hidden state $\mathbf{h}_t^s \in \mathbb{R}^K$ from the sharing layer, the previous hidden state \mathbf{h}_{t-1}^m from the USE-specific layer, and the current input \mathbf{x}_t^m . The hidden state \mathbf{h}_t^m of our GRU^m can be computed as below:

$$\begin{aligned} r_t^m &= \delta(\mathbf{W}_r^m \mathbf{x}_t^m + \mathbf{U}_r^m \mathbf{h}_{t-1}^m + \mathbf{U}_r^{sm} \mathbf{h}_t^s) \\ z_t^m &= \delta(\mathbf{W}_z^m \mathbf{x}_t^m + \mathbf{U}_z^m \mathbf{h}_{t-1}^m + \mathbf{U}_z^{sm} \mathbf{h}_t^s) \\ \hat{\mathbf{h}}_t^m &= \tau(\mathbf{W}_h^m \mathbf{x}_t^m + \mathbf{U}_h^m (\mathbf{h}_{t-1}^m \odot r_t^m) + \mathbf{U}_h^{sm} \mathbf{h}_t^s) \\ \mathbf{h}_t^m &= (1 - z_t^m) \odot \mathbf{h}_{t-1}^m + z_t^m \odot \hat{\mathbf{h}}_t^m \end{aligned} \quad (3)$$

where \mathbf{U}_*^{sm} denotes the weight matrix which connects the sharing layer and the task-specific layer, \mathbf{W}_*^m and \mathbf{U}_*^m are similar to the parameters of GRU^s , $\delta(\cdot)$ is the sigmoid function, $\tau(\cdot)$ is the hyperbolic tangent function and \odot is the hadamard product. For the *auxiliary* SA task, we also have the similar operation via the GRU^s from the sharing layer and a GRU^a from a SA-specific layer as below:

$$\begin{aligned} \mathbf{h}_t^s &= GRU^s(\mathbf{h}_{t-1}^s, \mathbf{x}_t^a) \\ \mathbf{h}_t^a &= GRU^a(\mathbf{h}_{t-1}^a, \mathbf{h}_t^s, \mathbf{x}_t^a) \end{aligned} \quad (4)$$

Decoder for User Satisfaction Estimation

User satisfaction is influenced by role and position information (Song et al. 2019; Liu et al. 2021b). Intuitively, user satisfaction is more related to the utterances of users rather than that of service staffs who usually express positive and neutral emotions, and utterances at the end of a dialogue tend to reflect the user’s final attitude. Thus, we first adopt a role-selected mask mechanism to reserve the hidden states of all the user utterances and then produce a user-specific vector. Besides, we measure the importance of each user utterance through a scoring function as below:

$$\begin{aligned} \mathbf{u}_t &= \tanh(\mathbf{W}_u \mathbf{h}_t^m + \mathbf{b}_u) \\ \alpha_t &= \frac{\exp(\text{mask}_r(\mathbf{u}_t^T \mathbf{U}_u))}{\sum_{k=1}^L \exp(\text{mask}_r(\mathbf{u}_k^T \mathbf{U}_u))} \end{aligned} \quad (5)$$

where $\mathbf{W}_u \in \mathbb{R}^{H \times K}$, $\mathbf{b}_u \in \mathbb{R}^H$ and $\mathbf{U}_u \in \mathbb{R}^H$ are trainable model parameters, function $mask_r(x)$ returns value x if speaker role $r = user$, otherwise $-\infty$. Therefore, we obtain the user-specific vector via the weighted sum of the hidden states of all the user utterances. In order to balance roles and positions, the final dialogue representation can be obtained by concatenating the user-specific vector and the last hidden output of GRU^m as follow:

$$\mathbf{o}^m = \left[\left(\sum_{t \in [1, L]} \alpha_t \mathbf{h}_t^m \right), \mathbf{h}_L^m \right] \quad (6)$$

where $[\cdot, \cdot]$ denotes the vector concatenation, the vector $\mathbf{o}^m \in \mathbb{R}^{2K}$ will be input into a linear layer and then a softmax layer for obtaining user satisfaction estimation as below:

$$\mathbf{p}^m = \text{softmax}(\mathbf{W}^m \mathbf{o}^m + \mathbf{b}^m) \quad (7)$$

where $\mathbf{W}^m \in \mathbb{R}^{|\Omega| \times 2K}$ and $\mathbf{b}^m \in \mathbb{R}^{|\Omega|}$ are learnable model parameters, and user satisfaction labels $\Omega = \{unsatisfied, met, well\ satisfied\}$.

Decoder for Sentiment Analysis

For the SA task, each utterance vector $\mathbf{h}_t^a \in \mathbb{R}^K$ from GRU^a unit will be first input into a linear layer and then a softmax layer for calculating the sentiment probability $\mathbf{p}_t^a \in \mathbb{R}^{|\mathcal{G}|}$:

$$\mathbf{p}_t^a = \text{softmax}(\mathbf{W}^a \mathbf{h}_t^a + \mathbf{b}^a) \quad (8)$$

where $\mathbf{W}^a \in \mathbb{R}^{|\mathcal{G}| \times K}$, $\mathbf{b}^a \in \mathbb{R}^{|\mathcal{G}|}$ are learnable parameters, $\mathcal{G} = \{negative, neutral, positive\}$ are sentiment classes.

The Training Procedure

For the USE and SA tasks, we have the single-task loss functions \mathcal{L}^{USE} and \mathcal{L}^{SA} , respectively, in Formula 9. To train our multi-task basic model, we sum the above loss functions up as the final multi-task loss function:

$$\mathcal{L}^{Task} = \underbrace{- \sum_{l \in \Omega} \mathbf{g}_l^m \log(\mathbf{p}_l^m)}_{\mathcal{L}^{USE} \text{ in main task}} - \frac{1}{L} \sum_{s \in \mathcal{D}} \sum_{l \in \mathcal{G}} \mathbf{g}_{sl}^a \log(\mathbf{p}_{sl}^a) \quad (9)$$

$\mathcal{L}^{SA} \text{ in auxiliary task}$

where \mathbf{g}_l^m and \mathbf{p}_l^m are respectively the ground truth and the predicted probability corresponding to the l -th satisfaction class. \mathbf{g}_{sl}^a and \mathbf{p}_{sl}^a are respectively the ground truth and the predicted probability corresponding to the l -th sentiment class for each utterance. We use back propagation to calculate the gradients of all the trainable parameters, and update them with momentum optimizer (Qian 1999).

STMAN: Speaker Turn-Aware Multi-Task Adversarial Network

In conventional multi-task learning, the performance of task models is influenced by the task-specific features. Intuitively, the dense layers in *Basic* model can not produce task-specific features well without supervision signals. Thus, we introduce additional *Multi-Task Adversarial Layer* which uses a *Task Discriminator* to differentiate task-specific input features explicitly. Besides, we consider speak-turn change information and introduce additional *Speaker Turn-Aware MT Interactive Layer*. Finally, we build a novel and extensible **Speaker Turn-Aware Multi-Task Adversarial Network (STMAN)** and display its architecture in Figure 3.

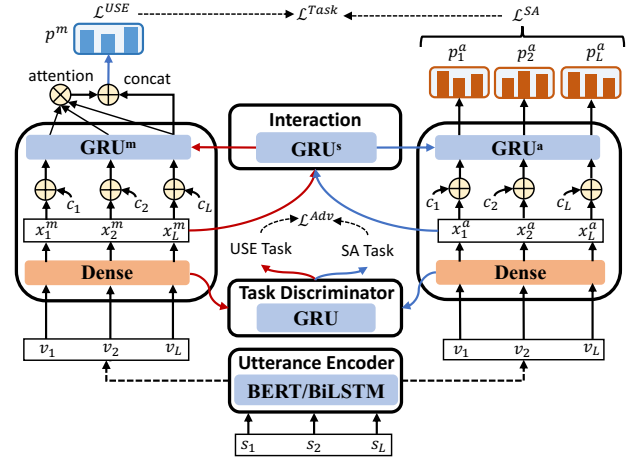


Figure 3: The architecture of our enhanced model: STMAN.

ST-Aware MT Interactive Layer

Different from speaker role-based methods (Qin et al. 2020), our method focuses on speaker turns and thus is still useful when speakers are not associated with specific roles. Specifically, we relabel the speakers and flip the speaker label (from 0 to 1 and vice versa) when there is speaker turn change. For the example in Figure 1, if the original speaker sequence is $\langle U, S, U, S, U, S, U, U, S \rangle^2$, the relabeling sequence is $\langle 0, 1, 0, 1, 0, 1, 0, 0, 1 \rangle$, which can be represented by the two introduced speaker turn embeddings in further. Given the utterance representation \mathbf{x}_t^* and its binary speaker turn label $r_{s_t} \in \{0, 1\}$, we have role enhanced utterance representation $[\mathbf{x}_t^*, \mathbf{c}_t]$, where $\mathbf{c}_t = Emb(r_{s_t})$ denotes any speaker turn embedding and $\mathbf{c}_t \in \mathbb{R}^Z$. Finally, the task-specific GRU units in Formulas 2 and 4 can be reformulated as below:

$$\begin{aligned} \mathbf{h}_t^m &= \text{GRU}^m(\mathbf{h}_{t-1}^m, \mathbf{h}_t^s, [\mathbf{x}_t^m, \mathbf{c}_t]) \\ \mathbf{h}_t^a &= \text{GRU}^a(\mathbf{h}_{t-1}^a, \mathbf{h}_t^s, [\mathbf{x}_t^a, \mathbf{c}_t]) \end{aligned} \quad (10)$$

Obviously, \mathbf{h}_t^m and \mathbf{h}_t^a contain speaker turn change information, which are helpful for downstream USE and SA tasks.

MT Adversarial Layer

Without any supervised signals, the two dense layers (see Formula 1) can not differentiate task-specific features accurately, thus the shared features can exist in the task private space and the task-specific features creep into the shared space. Intuitively, adversarial learning can be further applied to multi-task learning, thus task-specific features \mathbf{X}^m and \mathbf{X}^a will be more distinguishable. Here, we introduce a *Task Discriminator* (TD). Specifically, we exploit the min-max game between the shared utterance encoder and the task-specific dense layer with the task discriminator (Goodfellow et al. 2014). TD is composed of a GRU layer stacked with a softmax layer to estimate what kind of tasks the utterance representations come from, i.e., it maps the task-specific features into a probability distribution $\mathbf{p}^k \in \mathbb{R}^{|\mathcal{K}|}$, where task type $k \in \mathcal{K}$ and $\mathcal{K} = \{USE, SA\}$. Thus, we use a task

²U denotes a user and S represents a service staff.

Algorithm 1: Multi-Task Adversarial Learning

Input: Training set $\{\mathcal{D}\}$;**Initialized Parameters:** shared utterance encoder θ_{Enc} , task private decoders $\{\theta_k | k \in \mathcal{K}\}$ and task discriminator θ_{TD} ;

- 1: **for** mini-batches in $\{\mathcal{D}\}$ **do**
- 2: Minimize \mathcal{L}^{Adv} in Eq. 11 to update θ_{Enc} ;
- 3: Maximize \mathcal{L}^{Adv} in Eq. 11 to update θ_{TD} and $\{\theta_k\}$;
- 4: **end for**
- 5: **for** mini-batches in $\{\mathcal{D}\}$ **do**
- 6: Minimize \mathcal{L}^{Task} in Eq. 9 to update all the trainable parameters except for θ_{TD} ;
- 7: **end for**

Output: The well trained θ_{Enc} , θ_{TD} and $\{\theta_k\}$

Statistics items	Clothes	Makeup
# Dialogues	10,000	3,540
# US (unsatisfied)	2,302	1,180
# MT (met)	6,399	1,180
# WS (well satisfied)	1,299	1,180
# Utterances	123,242	46,255
# NG (negative)	12,619	6,130
# NE (neutral)	97,380	33,158
# PO (positive)	13,243	6,976

Table 1: The statistics of the two benchmark datasets.

adversarial loss \mathcal{L}^{Adv} to limit the task-specific features into their private feature space, which is defined as follow:

$$\mathcal{L}^{Adv} = \min_{\theta_{Enc}} \left(\max_{\theta_{TD}, \theta_k} \sum_{s \in \mathcal{D}} \sum_{l \in \mathcal{K}} \mathbf{g}_{sl}^k \log(\mathbf{p}_{sl}^k) \right) \quad (11)$$

where \mathbf{g}_{sl}^k and \mathbf{p}_{sl}^k are respectively the ground-truth and the predicted probability corresponding to the l -th task class for each utterance in the dialogue \mathcal{D} . θ_{Enc} , θ_{TD} and $\{\theta_k\}$ are the learnable parameters of *Shared Utterance Encoder*, *Task Discriminator* and respective *Decoders*. And the whole training procedure is shown in algorithm 1.

Experiments and Results

In this section, we conduct extensive experiments on two publicly available service dialogue datasets to evaluate the effectiveness of our approach. *Clothes* and *Makeup* datasets are collected by Song et al. (2019) from E-commerce platform taobao.com and split into training/development/testing sets with splits 8/1/1. The *Clothes* is a larger imbalanced dataset and the *Makeup* is a smaller balanced dataset. The statistics of datasets are given in Table 1.

Experimental Settings

We encode utterances with BiLSTM as that used in works (Song et al. 2019; Liu et al. 2021b) because BERT-based encoder is a sub-optimal choice considering different types of pre-training corpus. Trainable model parameters are initialized by sampling values from a uniform distribution $\mathcal{U}(-0.01, 0.01)$. The hyper-parameters are tuned to the best on the development set. The size of hidden states K is 100, speaker turn embedding size Z is 100, H is 50, the dropout

rate is 0.2, the learning rate is 0.1, the learning rate decay is 0.8, the batch size is 32 and the number of epochs is 30. We use Macro F1 and Accuracy as the evaluation metrics. Our programs are implemented by tensorflow³ and run on a server configured with a Tesla V100 GPU, 2 CPU and 32G memory. All the resources are publicly available⁴.

Comparative Study

In Table 2 and Table 3, we compare our approach with several state-of-the-art USE single-task models (1-4) and the multi-task models (5-8) over the two dialogue datasets.

- **BERT+LSTM** encodes each utterance with BERT-based encoder and uses the last hidden state of LSTM for satisfaction estimation (Liu et al. 2021b).
- **MILNET** represents a Multiple Instance Learning Network for document-level and sentence-level sentiment analysis (Angelidis and Lapata 2018). In Song et al. (2019), USE is considered as a special SA task by treating dialogue as document and utterance as sentence.
- **CAMIL** is a Context-Assisted Multiple Instance Learning model which predicts the utterance-level sentiments of all the user utterances and then aggregates the sentiments into user satisfaction estimation (Song et al. 2019).
- **SLUP** denotes a session-level user satisfaction estimation method trained on a sequence of question-answer pairs. To adapt to this method, we produce Q-A pairs by considering consecutive user utterances as questions and consecutive staff utterances as answers (Yao et al. 2020).
- **MT-ES** is an Enhanced Shared-layer Multi-Task architecture with two hidden layers: one is used to extract the common patterns via the shared parameters, and the other is used to capture task-specific features via the separate parameter set (Ma, Gao, and Wong 2018).
- **MT-US** is a simplified version of MT-ES with the uniform shared-layer (Cerisara et al. 2018). This belongs to a kind of shared-bottom multi-task model.
- **Meta-LSTM** is a shared meta-network to capture the meta-knowledge of semantic composition, which conducts text classification and sequence tagging (Chen et al. 2018). We input the utterance vectors directly, and output dialogue-level USE label and utterance-level SA labels.
- **RSSN** conducts machine-human chatting handoff (i.e., *transferable* or *normal*) and service satisfaction analysis jointly. For the three-class SA task, RSSN is tailored and tuned best on the development sets (Liu et al. 2021b).

Note that MILNET and CAMIL output utterances' sentiment predictions implicitly derived from satisfaction labels.

Results and Analysis: From Table 2 and Table 3, we can find MILNET performs worse than SLUP because it ignores dialogue structure information. SLUP achieves substantial improvements by modeling dialogue interactions precisely, but it still ignores the usage of sentiment knowledge. CAMIL performs better than BERT+LSTM by aggregating sentiment predictions directly into satisfaction polarity, however its cascade structure may introduce noisy

³<https://tensorflow.google.cn/>⁴<https://sites.google.com/view/ssa-sa/>

Methods	USE (main)		SA (auxiliary)	
	Macro F1	Accuracy	Macro F1	Accuracy
BERT+LSTM (Liu et al. 2021b)	67.90	75.50	-	-
MILNET (Song et al. 2019)	63.81	75.34	55.32	71.33
SLUP (Yao et al. 2020)	66.56	75.70	-	-
CAMIL (Song et al. 2019)	70.40	78.31	64.44	82.43
MT-US (Cerisara et al. 2018)	68.25	74.60	80.10	92.65
MT-ES (Majumder et al. 2019)	68.04	74.20	80.19	90.78
Meta-LSTM (Chen et al. 2018)	69.05	77.40	80.08	92.56
RSSN (Liu et al. 2021b)	69.51	77.20	80.29	91.15
STMAN (See Figure 3)	71.11	78.60	81.90	93.05

Table 2: Comparison among different models on *Clothes*.

Methods	USE (main)		SA (auxiliary)	
	Macro F1	Accuracy	Macro F1	Accuracy
BERT+LSTM (Liu et al. 2021b)	74.80	74.80	-	-
MILNET (Song et al. 2019)	75.30	75.19	41.71	41.09
SLUP (Yao et al. 2020)	76.31	76.28	-	-
CAMIL (Song et al. 2019)	78.60	78.59	59.54	64.73
MT-US (Cerisara et al. 2018)	76.37	75.98	82.12	90.77
MT-ES (Majumder et al. 2019)	76.06	76.83	80.51	89.96
Meta-LSTM (Chen et al. 2018)	76.65	76.55	80.13	90.48
RSSN (Liu et al. 2021b)	79.18	79.17	83.36	90.42
STMAN (See Figure 3)	80.11	80.22	84.18	91.54

Table 3: Comparison among different models on *Makeup*.

features. MT-US, MT-ES and Meta-LSTM are typical MT models, which perform not as well as RSSN because they are general frameworks and ignore dialogue structure which influences model performance greatly. Compared with all these baselines, our STMAN achieves the best results in both the tasks because it obtains better task-specific features by MT adversarial strategy and models utterance interaction well by ST-aware MT interaction strategy.

Ablation Study

Different model configurations can largely affect the model performance. In Table 4, we implement several model variants for ablation study by removing (“-”) and adding (“+”) different model components. **Basic** is the basic method ignoring all the strategies (see Figure 2). **Basic-Mask** removes mask operation in Formula 5 and reserves all the speakers’ utterances. **Basic-Aux** removes auxiliary SA task. **Basic+TD** enhances **Basic** by considering task discriminator. **Basic+ST** enhances **Basic** by considering speaker turn changes. **Basic+TD+ST** is our fully configured model.

From Table 4, we can find that **Basic** model performs better than **Basic-Mask** and **Basic-Aux**, which validates the usefulness of considering user utterances and speaker sentiments. **Basic+ST** perform better than **Basic** because speaker turn change provides helpful clues for utterance representation. Although TD is a simple classifier and its capability can be limited, **Basic+TD** still contributes stable and consistent performance improvements. The performance of **Basic+TD+ST** performs best among all these methods, which implies the effectiveness of different components.

Methods on Clothes	USE (main)		SA (auxiliary)	
	Macro F1	Accuracy	Macro F1	Accuracy
Basic (See Figure 2)	67.67	74.40	81.05	92.36
Basic-Mask	67.29	73.70	79.70	92.26
Basic-Aux	67.03	73.31	-	-
Basic+TD	68.92	75.80	81.52	92.64
Basic+ST	69.53	76.70	81.54	92.48
Basic+TD+ST (STMAN)	71.11	78.60	81.90	93.05

Methods on Makeup	USE (main)		SA (auxiliary)	
	Macro F1	Accuracy	Macro F1	Accuracy
Basic (See Figure 2)	78.69	78.71	83.01	88.96
Basic-Mask	77.12	77.11	82.41	88.62
Basic-Aux	74.64	74.57	-	-
Basic+TD	79.05	79.09	83.03	91.30
Basic+ST	79.03	79.05	83.30	91.22
Basic+TD+ST (STMAN)	80.11	80.22	84.18	91.54

Table 4: Comparison among different model configurations.




Figure 4: The influences of TD on USE and SA tasks.

Influences of Adversarial Learning

To investigate the influences of our multi-task adversarial learning strategy, we perform a in-depth study on *Task Discriminator* (TD). Specifically, we randomly sample $x\%$ ($x = 0, 25, 50, 75, 100$) training instances from the original training set, and further study the testing performance of USE and SA tasks, respectively. The experimental results are displayed in Figure 4.

From Figure 4, we can find that our method performs better as the size of training data increases. Our method performs worst when $x = 0$ because TD is not trained and it can not differentiate between USE and SA tasks. Our method performs quite well when $x \geq 75$ because TD achieves stable classification performance, which indicates that input features are more task-specific.

Correlation between USE and SA

To gain further insight into the correlation between USE and SA tasks, we enumerate all the (initial/final user sentiment, user satisfaction) combinations and calculate its proportion over the two datasets. Note that the initial/final user sentiment corresponds to the first/last user utterance in any dialogue. Figure 5 displays the statistical results, from which we can find some interesting patterns: (i) Initial/final neutral sentiment will trigger any user satisfaction class, which is more intuitive on the *Makeup* dataset. (ii) Initial negative sentiment is not effectively alleviated, but it gets worse at

Initial Sentiment in Makeup				Final Sentiment in Makeup			
PO	0.002	0.257	0.074	0.001	0.105	0.227	
NE	0.003	0.304	0.027	0.001	0.266	0.067	
NG	0.072	0.231	0.030	0.153	0.155	0.025	
	US	MT	WS	US	MT	WS	
Initial Sentiment in Clothes				Final Sentiment in Clothes			
PO	0.001	0.094	0.035	0.000	0.039	0.090	
NE	0.009	0.568	0.063	0.013	0.518	0.110	
NG	0.087	0.126	0.018	0.129	0.087	0.014	
	US	MT	WS	US	MT	WS	

Figure 5: The proportion of all the (initial/final user sentiment, user satisfaction) combinations over the datasets.

Methods	Clothes		Makeup	
	Macro F1	Accuracy	Macro F1	Accuracy
Initial-User-Sentiment	53.83	68.60	43.23	47.74
Final-User-Sentiment	65.96	72.30	64.51	64.40
STMAN	71.11	78.60	80.11	80.22

Table 5: Initial/final user sentiments versus user satisfaction.

the end. Because the proportion of (US, NG) increases from 7.2% to 15.3% on *Makeup* and 8.7% to 12.9% on *Clothes*, respectively. (iii) Compared with initial sentiment, final sentiment plays more critical roles on deciding the user satisfaction. Specifically, we directly map the initial/final negative, neutral and positive sentiment to the corresponding user satisfaction, i.e., unsatisfied, met and satisfied, respectively. Given final sentiment, the total proportion of (US, NG), (MT, NE) and (WS, PO) combinations is significantly larger than under initial sentiments. All these verify that USE and SA are correlated tasks and can be learned together.

Besides, we follow the heuristic *sentiment-satisfaction* mapping method introduced above, and display the experimental results in Table 5. From the results, we can find that although SA and USE are highly pertinent tasks, sentiments can not be mapped to the user satisfaction directly, especially for *met* satisfaction which has less connection with sentiments. This again verifies that SA task only provides limited helpful sentimental clues to assist with USE.

Case Study

To better understand the usefulness of our model, we choose three example dialogues from *Clothes* testing set and display the prediction results of Meta-LSTM, RSSN and STMAN in Table 6. The original Chinese dialogues have been translated into English dialogues for understandability. We denote user as “U” and service staff as “S” to differentiate speakers.

For the 1st example, the user expresses negative sentiment in the beginning, but he is well comforted and expresses neutral sentiment in the end after the service staff explains the

Exemplar Dialogues (Chinese → English)	
U: Can you speed up logistics? Too slow! (NG,NG,NG,NG)	
S: The logistics is running normally. (NE,NE,NE,NE)	
——Omit staff explanation——	
U: Try to help me speed up the progress. (NG,NG,NE,NE)	
S: Dear, I’ll try my best to give feedback. (PO,PO,PO,PO)	
S: Thanks for your support! Bye! (PO,PO,PO,PO)	
Meta-LSTM:US RSSN:US STMAN:MT Truth:MT	
U: Hello! (NE,NE,NE,NE)	
S: What can I do for you, dear? (PO,PO,PO,PO)	
U: What about the back of the pants?	
Do you have any pictures? (NE,NE,NE,NE)	
S: Dear, there is no real picture here. (PO,PO,PO,PO)	
S: Dear, please check other people’s reviews. (PO,PO,PO,PO)	
U: If you don’t know your product, how can you introduce it to your customers? ? (NE,NE,NG,NG)	
S: Not every item has pictures. :):) (PO,PO,PO,PO)	
Meta-LSTM:MT RSSN:US STMAN:US Truth:US	
U: I don’t want ZTO Express. (NE,NE,NE,NE)	
U: Are you there? (NE,NE,NE,NE)	
S: Please wait a moment. (NE,NE,NE,NE)	
——No reply for a long time——	
U: ? ? ? (NE,NE,NE,NE)	
Meta-LSTM:MT RSSN:MT STMAN:MT Truth:US	

Table 6: The predicted results of three example dialogues in *Clothes*. The sentiment and satisfaction are displayed in Quadruple (Meta-LSTM, RSSN, STMAN, Truth).

reason very sincerely (*detail content is omitted*). However, Meta-LSTM and RSSN mispredict the 3rd utterance which is similar to the 1st utterance in content. For the 2nd example, Meta-LSTM and RSSN can not predict the correct sentiment of 6th utterance which implies accusation. For the 3rd example, all the three methods make wrong predictions because it lacks of useful information and its sentiment is dependent on time factor which is not considered in modeling. For the hard-to-solve case, we leave it for future study.

Conclusion and Future Work

In this paper, we propose a novel and extensible STMAN for joint USE and SA. Firstly, we introduce a basic model which adopts two hidden layers for each task: one is used to extract the common patterns via the shared parameters, and the other is used to capture task-specific features via the separate parameter set. Then, we enhance the basic model by introducing speak turn information to enhance the interaction layer. Finally, we introduce a multi-task adversarial strategy which makes inputs more task-specific via a task discriminator. Experiments conducted on two public service dialogue datasets indicate the effectiveness of our approach.

In the future, we will investigate how to estimate user satisfaction in low-resource scenario.

Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions. This work is supported by National Natural Science Foundation of China (62106039) and National Key R&D Program of China (2020YFC0832505).

References

- Abraham, M. M.; and Choe, P. 2021. Factors Affecting e-Commerce Satisfaction in Qatar: A Cross-Cultural Comparison. In *Proceedings of HCI 2021*, 481–494.
- Angelidis, S.; and Lapata, M. 2018. Multiple Instance Learning Networks for Fine-Grained Sentiment Analysis. *TACL*, 6: 17–31.
- Bodigutla, P. K.; Tiwari, A.; Matsoukas, S.; Valls-Vargas, J.; and Polymenakos, L. 2020. Joint Turn and Dialogue level User Satisfaction Estimation on Multi-Domain Conversations. In *EMNLP (Findings) 2020*, 3897–3909.
- Caruana, R. 1997. Multitask Learning. *Mach. Learn.*, 28(1): 41–75.
- Cerisara, C.; Jafaritazehjani, S.; Oluokun, A.; and Le, H. T. 2018. Multi-task dialog act and sentiment recognition on Mastodon. In *Proceedings of the 27th COLING*, 745–754.
- Chen, J.; Qiu, X.; Liu, P.; and Huang, X. 2018. Meta Multi-Task Learning for Sequence Modeling. In *Proceedings of the 32nd AAAI*, 5070–5077.
- Cho, K.; van Merriënboer, B.; Gülçehre, Ç.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of EMNLP 2014*, 1724–1734.
- Deng, Y.; Zhang, W.; Lam, W.; Cheng, H.; and Meng, H. 2022. User Satisfaction Estimation with Sequential Dialogue Act Modeling in Goal-oriented Conversational Systems. In *Proceedings of the 31st WWW*, 2998–3008.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT 2019*, 4171–4186. ACL.
- Gadri, S.; Mehieddine, S. O.; Herizi, K.; and Chabira, S. 2021. An Efficient System to Predict Customers’ Satisfaction on Touristic Services Using ML and DL Approaches. In *Proceedings of the 22nd ACIT*, 1–8.
- Gao, J.; and Zhang, Z. 2011. User Satisfaction of Ali Wangwang, an Instant Messenger Tool. In *Proceedings of the 14th HCI*, volume 6770, 414–420.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A. C.; and Bengio, Y. 2014. Generative Adversarial Nets. In *Proceedings of NIPS 2014*, 2672–2680.
- He, Z.; Tavabi, L.; Lerman, K.; and Soleymani, M. 2021. Speaker Turn Modeling for Dialogue Act Classification. In *Proceedings of EMNLP (Findings) 2021*, 2150–2157.
- Higashinaka, R.; D’Haro, L. F.; Shawar, B. A.; Banchs, R. E.; Funakoshi, K.; Inaba, M.; Tsunomori, Y.; Takahashi, T.; and Sedoc, J. 2019. Overview of the Dialogue Breakdown Detection Challenge 4. In *Proceedings of IWSIDS 2019*, 403–417.
- Kachuee, M.; Yuan, H.; Kim, Y.; and Lee, S. 2021. Self-Supervised Contrastive Learning for Efficient User Satisfaction Prediction in Conversational Agents. In *Proceedings of NAACL-HLT 2021*, 4053–4064.
- Liu, J.; Gao, Z.; Kang, Y.; Jiang, Z.; He, G.; Sun, C.; Liu, X.; and Lu, W. 2021a. Time to transfer: Predicting and evaluating machine-human chatting handoff. In *Proceedings of the 35th AAAI*, 5841–5849.
- Liu, J.; Song, K.; Kang, Y.; He, G.; Jiang, Z.; Sun, C.; Lu, W.; and Liu, X. 2021b. A Role-Selected Sharing Network for Joint Machine-Human Chatting Handoff and Service Satisfaction Analysis. In *Proceedings of EMNLP 2021*, 9731–9741.
- Ma, J.; Gao, W.; and Wong, K. 2018. Detect Rumor and Stance Jointly by Neural Multi-task Learning. In *Proceedings of the 27th WWW*, 585–593.
- Majumder, N.; Poria, S.; Hazarika, D.; Mihalcea, R.; Gelbukh, A. F.; and Cambria, E. 2019. DialogueRNN: An Attentive RNN for Emotion Detection in Conversations. In *Proceedings of the 33rd AAAI*, 6818–6825.
- Mehrotra, R.; Lalmas, M.; Kenney, D.; Lim-Meng, T.; and Hashemian, G. 2019. Jointly Leveraging Intent and Interaction Signals to Predict User Satisfaction with Slate Recommendations. In *Proceedings of the 28th WWW*, 1256–1267.
- Qian, N. 1999. On the momentum term in gradient descent learning algorithms. *Neural Networks*, 12(1): 145–151.
- Qin, L.; Che, W.; Li, Y.; Ni, M.; and Liu, T. 2020. DCR-Net: A Deep Co-Interactive Relation Network for Joint Dialog Act Recognition and Sentiment Classification. In *Proceedings of the 34th AAAI*, 8665–8672.
- Qin, L.; Li, Z.; Che, W.; Ni, M.; and Liu, T. 2021. Co-GAT: A Co-Interactive Graph Attention Network for Joint Dialog Act Recognition and Sentiment Classification. In *Proceedings of the 35th AAAI*, 13709–13717.
- Schuster, M.; and K.Paliwal, K. 1997. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.*, 45(11): 2673–2681.
- Shan, Y.; Cui, A.; Tan, L.; and Xiong, K. 2019. Overview of the NLPCC 2019 Shared Task: Open Domain Conversation Evaluation. In *Proceedings of NLPCC 2019*, 829–834.
- Song, K.; Bing, L.; Gao, W.; Lin, J.; Zhao, L.; Wang, J.; Sun, C.; Liu, X.; and Zhang, Q. 2019. Using Customer Service Dialogues for Satisfaction Analysis with Context-Assisted Multiple Instance Learning. In *Proceedings of EMNLP-IJCNLP 2019*, 198–207.
- Song, S.; Wang, C.; Chen, H.; and Chen, H. 2021. An Emotional Comfort Framework for Improving User Satisfaction in E-Commerce Customer Service Chatbots. In *Proceedings of NAACL 2021*, 130–137.
- Song, S.; Wang, C.; Liu, S.; Chen, H.; Chen, H.; and Bao, H. 2020. Sentiment Analysis Technologies in AliMe - An Intelligent Assistant for E-commerce. *Int. J. Asian Lang. Process.*, 30(4): 2050016:1–2050016:20.
- Sun, W.; Zhang, S.; Balog, K.; Ren, Z.; Ren, P.; Chen, Z.; and de Rijke, M. 2021. Simulating User Satisfaction for the Evaluation of Task-oriented Dialogue Systems. In *Proceedings of the 44th SIGIR*.
- Wang, J.; Wang, J.; Sun, C.; Li, S.; Liu, X.; Si, L.; Zhang, M.; and Zhou, G. 2020. Sentiment Classification in Customer Service Dialogue with Topic-Aware Multi-Task Learning. In *Proceedings of the 34th AAAI*, 9177–9184.

Yao, R.; Song, S.; Li, Q.; Wang, C.; Chen, H.; Chen, H.; and Zeng, D. D. 2020. Session-Level User Satisfaction Prediction for Customer Service Chatbot in E-Commerce (Student Abstract). In *Proceedings of the 34th AAAI*, 13973–13974.

Zhou, B.; Cai, X.; Zhang, Y.; Guo, W.; and Yuan, X. 2021. MTAAL: Multi-Task Adversarial Active Learning for Medical Named Entity Recognition and Normalization. In *Proceedings of the 35th AAAI*, 14586–14593.