

# Self-Supervised Contrastive Learning for Efficient User Satisfaction Prediction in Conversational Agents

**Mohammad Kachuee\***      **Hao Yuan**      **Young-Bum Kim**      **Sungjin Lee**  
UCLA Computer Science    Amazon Alexa AI      Amazon Alexa AI      Amazon Alexa AI  
Los Angeles, CA          Seattle, WA            Seattle, WA            Seattle, WA  
mkachuee@cs.ucla.edu      {yuanha, youngbum, sungjinl}@amazon.com

## Abstract

Turn-level user satisfaction is one of the most important performance metrics for conversational agents. It can be used to monitor the agent’s performance and provide insights about defective user experiences. While end-to-end deep learning has shown promising results, having access to a large number of reliable annotated samples required by these methods remains challenging. In a large-scale conversational system, there is a growing number of newly developed skills, making the traditional data collection, annotation, and modeling process impractical due to the required annotation costs and the turnaround times. In this paper, we suggest a self-supervised contrastive learning approach that leverages the pool of unlabeled data to learn user-agent interactions. We show that the pre-trained models using the self-supervised objective are transferable to the user satisfaction prediction. In addition, we propose a novel few-shot transfer learning approach that ensures better transferability for very small sample sizes. The suggested few-shot method does not require any inner loop optimization process and is scalable to very large datasets and complex models. Based on our experiments using real data from a large-scale commercial system, the suggested approach is able to significantly reduce the required number of annotations, while improving the generalization on unseen skills.

## 1 Introduction

Nowadays automated conversational agents such as Alexa, Siri, Google Assistant, Cortana, etc. are widespread and play an important role in many different aspects of our lives. Their applications vary from storytelling and education for children to assisting the elderly and disabled with their daily activities. Any successful conversational agent should be able to communicate in different languages and accents, understand the conversation

context, analyze the query paraphrases, and route the requests to various skills available for handling the user’s request (Ram et al., 2018).

In such a large-scale system with many components, it is crucial to understand if the human user is satisfied with the automated agent’s response and actions. In other words, it is desirable to know if the agent is communicating properly and providing the service that is expected by the user. In the literature, it is referred to as targeted turn-level satisfaction as we are only interested in the user’s satisfaction for a certain conversation turn given the context of the conversation, and not the overall satisfaction for the whole conversation (Park et al., 2020). Perhaps the most basic use of a user satisfaction model would be to monitor the performance of an agent and to detect defects as a first step to fix issues and improve the system. Anticipating user dissatisfaction for a certain turn in a conversation, an agent would be able to ask the user for repeating the request or providing more information, improving the final experience. Also, a powerful user satisfaction model can be used as a ranking or scoring measure to select the most satisfying response among a set of candidates and hence guiding the conversation.

The problem of user satisfaction modeling has recently attracted significant research attention (Jiang et al., 2015; Bodigutla et al., 2019; Park et al., 2020; Pragst et al., 2017; Rach et al., 2017). These methods either rely on annotated datasets providing ground-truth labels to train and evaluate (Bodigutla et al., 2019) or rely on ad hoc or human-engineered metrics that do not necessarily model the true user satisfaction (Jiang et al., 2015). Access to reliable annotations to be used in building satisfaction models has been very challenging partly due to the fact that a large-scale conversation system supports many different devices as well as voice, language, and application components, providing access to a wide variety of skills. The traditional approach of collecting samples from the live system traffic and

\*Work done as an intern at Amazon Alexa AI.

tasking human annotators to label samples would not be scalable due to the cost of annotations as well as the turn-around time required to collect and annotate data for a new skill or feature. Note that onboarding new skills in a timely manner is a crucial to ensure active skill developer engagement.

To address this problem, we propose a novel training objective and transfer learning scheme that significantly improves not only the data efficiency but also the model generalization to unseen skills. In summary, we make the following contributions:

- We propose a contrastive self-supervised training objective that can leverage virtually any unlabeled conversation data to learn user-agent interactions.
- We show that the proposed method can be used to pre-train state-of-the-art deep language models and the acquired knowledge is transferable to the user satisfaction prediction.
- We suggest a novel and scalable few-shot transfer learning approach that is able to improve the label efficiency even further in the case of few-shot transfer learning.
- We conduct extensive experiments using data from a large-scale commercial conversational system, demonstrating significant improvements to label efficiency and generalization.

## 2 Related Work

### User Satisfaction in Conversational Systems

The traditional approach to evaluating a conversational system is to evaluate different functionalities or skills individually. For instance, for a knowledge question answering or web search skill, one can use response quality metrics commonly used to evaluate search system and ranking systems such as nDCG (Järvelin et al., 2008; Hassan, 2012; Fox et al., 2005). While these methods provide justifiable measures for certain skills, they are not extendable to a large number of skills, especially for skills without a set of proper hand-engineered features and metrics, or newly developed third-party skills (Bodigutla et al., 2019).

Another, more general, line of research is to evaluate the performance of a conversation system from the language point of view. Here, the objective is to measure how natural, syntactically and semantically, an automated agent is able to

interact with a human user. For instance, using generic metrics such as BLEU (Papineni et al., 2002) or ROUGE (Lin, 2004) one can measure how the agent’s responses are consistent with a set of provided ground-truth answers. However, these approaches not only suffer from shortcomings such as inconsistency with the human understanding (Liu et al., 2016; Novikova et al., 2017) but also are not practical for a real-world conversation system due to their dependence on ground-truth responses.

A more recent approach is to use human annotations specifically tailored for the user satisfaction task as a source of supervision to train end-to-end prediction models (Bodigutla et al., 2019). Jiang et al. (2015) suggested training individual models for 6 general skills and devised engineered features to link user actions to the user satisfaction for each studied skill. Park et al. (2020) proposed a hybrid method to learn from human annotation and user feedback data that is scalable and able to model user satisfaction across a large number of skills.

### Contrastive Learning

Gutmann and Hyvärinen (2010) was the first study to propose the idea of noise-contrastive learning in the context of a capturing a distribution using an objective function to distinguish samples of the target distribution from samples of an artificially generated noise distribution. Contrastive predictive coding (CPC) (Oord et al., 2018) suggested the idea of using an NCE objective to train an autoregressive sequence representation model. Deep InfoMax (Hjelm et al., 2018) used self-supervised contrastive learning in an architecture where a discriminator is trained to distinguish between representations of the same image (positive samples) or representations of different images (negative samples). While many different variations of contrastive methods have been suggested, the main idea remains the same: defining a self-supervised objective to distinguish between the hidden representations of samples from the original distribution and samples from a noise distribution (Trinh et al., 2019; Devon et al., 2020; Yao et al., 2020).

### Few-shot Transfer Learning

Few-shot transfer learning is a very active and broad subject of research. We limit the scope of our study to methods in which a form of gradient supervision is provided by a target task to ensure the efficient transferability of representations trained on a source task. Lopez-Paz and Ranzato (2017)

suggested the idea of joint multi-task training and using the cosine similarity of the concatenated network gradients from the source and target tasks. For gradients with negative cosine distance, they project the source gradients to a more aligned direction by solving a quadratic programming problem. Luo et al. (2020) continued that line and suggested a method in the context of few-shot transfer learning, showing that using even a few samples from the target task can significantly improve the transferability of the trained models. Li et al. (2020) presented a similar idea but suggested adjusting learning rates for each layer to improve the cosine similarity of different tasks. While these methods show promising results, they only measure the similarity between concatenated gradient vectors consisting of all network parameters which is a very rough measure of alignment. Also, they require solving for a quadratic or iterative optimization problem as an inner loop in the training procedure that can be computationally expensive and often prohibitive for large-scale problems.

### 3 User Satisfaction Modeling

#### 3.1 Problem Definition

In this paper, we consider the conversational interaction between a human user and an automated agent. Each interaction consists of a set of turns in which the user provides an utterance and the agent provides appropriate responses. A set of turns that are happening within a certain time window are grouped as a conversation session. Formally, we can represent a session as a set of turns:

$$S_i = \{(U_i^{t=0}, R_i^{t=0}), \dots, (U_i^{t=T}, R_i^{t=T})\} \quad (1)$$

Here,  $S_i$  represents session  $i$  consisting of a set of turns as tuples of utterance and responses,  $(U_i^t, R_i^t)$ , for the first turn  $t = 0$  to the last turn  $t = T$  in that session.

In the context of turn-level user satisfaction modeling, we are interested in the classification of a certain targeted turn within a session as either satisfying (SAT) or dissatisfying (DSAT). Note that the satisfaction here is defined based on the agent’s response given a certain utterance and the context (i.e., other session turns). We use the notation  $Y_i^{t^*} \in \{\text{SAT}, \text{DSAT}\}$  to indicate the user satisfaction for the targeted turn  $t = t^*$  of session  $i$ . See Figure 2 for examples of SAT/DSAT interactions.

Session	Label
<b>U0:</b> <i>play clocks</i> <b>R0:</b> <i>Playing clocks by Coldplay.</i> <i>[User listening for 1 minute]</i>	<b>SAT</b> (user listening)
<b>U0:</b> <i>play clock</i> <b>R0:</b> <i>Sorry, I cannot find the song.</i> U1: <i>play clocks</i> R1: <i>Playing clocks by Coldplay.</i> <i>[User listening for 1 minute]</i>	<b>DSAT</b> (repeated request)
<b>U0:</b> <i>turn on the tv</i> <b>R0:</b> <i>Did you mean Living Room TV?</i> U1: <i>yes</i> R1: <i>[EMPTY]</i>	<b>SAT</b> (no further request)
<b>U0:</b> <i>call chris</i> <b>R0:</b> <i>Calling chris.</i> U1: <i>stop</i> R1: <i>[EMPTY]</i> U1: <i>call christina</i> R1: <i>Calling Christina.</i> <i>[Phone call for 2 minutes]</i>	<b>DSAT</b> (ASR error, barge-in)

Figure 2: A few examples of SAT and DSAT turns to illustrate the importance of the conversation context.

#### 3.2 Datasets<sup>1</sup>

In this study, we use real-world data from Alexa, a large-scale commercial conversational agent. Specifically, we use a dataset of about 891,000 real-world conversation sessions in which a certain turn within each session is annotated by a human annotator as SAT or DSAT. Human annotators had access to the session context and followed a standard labeling protocol (further information is provided in Appendix A). As a preprocessing step, we limited turns within each session to a window of five turns: at most two turns before the targeted turn, the targeted turn, and at most two turns after the targeted turn. This labeled dataset is denoted as  $\mathbb{D}_{sup}$ .

In addition to  $\mathbb{D}_{sup}$ , we also use a large pool of real-world session data without any annotation or label. This dataset is about twice the size of  $\mathbb{D}_{sup}$ , but as we are not limited to targeted turns, we keep all session turns and decide context windows based on a randomized data augmentation step. The resulting effective sample size is significantly larger than  $\mathbb{D}_{sup}$ . We denote this unlabeled dataset as  $\mathbb{D}_{unsup}$ . As both datasets were sampled from real traffic, we ensured that there is no overlap between  $\mathbb{D}_{unsup}$  and the evaluation splits of  $\mathbb{D}_{sup}$ .

The conversations cover a wide variety of internally developed (1p) and third-party (3p) developer skills. Due to the imbalanced traffic, in our

<sup>1</sup>Due to confidentiality concerns, we are not able to disclose the exact annotation protocols and data specifications.




Figure 1: Overview of the suggested network architecture. In our architecture, BERT encoder with average pool at the last layer is used as the LM. We consider a context window of at most  $2T+1$  turns. Heads are simple MLPs classifiers with one hidden layer.

Property	Size
Total number of samples	$\approx 891,000$
Total number of 1p skills	$> 20$
Total number of 3p skills	$> 1500$
Ratio of SAT to DSAT samples	$> 20$

Table 1: Dataset statistics for  $\mathbb{D}_{sup}$

datasets, there is a huge variation between the number of samples for different skills. For instance, 1p skills such as music or weather have hundreds of thousands of samples while many 3p skills only have less than 10 samples throughout our datasets. To properly evaluate the performance of our predictors on such imbalanced data, we proposed a novel approach to split the data and to evaluate. We build two test sets: a test set measuring in-domain performance and another test set to measure the out-of-domain generalization. The in-domain test set consists of samples from skills that the train set covers. The out-of-domain test set measures the performance on skills that are not covered by the train set. Ideally, we would like to observe good classification performance in both test splits, indicating the ability of our models to learn and model the current major traffic and to generalize to less frequent or future traffic. Based on this, we split  $\mathbb{D}_{sup}$  to 70% train, 15% validation, and the rest for the test (about  $1/5$  of test samples are out-of-domain and  $4/5$  are in-domain). The in-domain and out-of-domain test sets consist of 17 and 275 skills, respectively. The  $\mathbb{D}_{unsup}$  is randomly split to 80% train and the rest for validation, regardless of skills. Table 1 presents a summary of dataset statistics for  $\mathbb{D}_{sup}$ .

### 3.3 Network Architecture

Figure 1 shows a high-level drawing of the network architecture used in our experiments. It consists of

a language model (LM) that encodes utterance and response pairs to vector representations. Here, we consider up to  $T$  turns before and after the targeted turn. To further summarize the list of the previous or next turns, we use GRU layers (Chung et al., 2014). Then, an average pool is used to produce a representation vector,  $z$ , for each session. Note that before the pooling, simple non-linear MLPs are used to transform each partial representation. Finally,  $z$  is used as an input to a set of different head networks, responsible for making predictions for different objectives.

Regarding the LM, we use the standard BERT encoder (Devlin et al., 2018) architecture pre-trained as suggested by Liu et al. (2019). To make a fixed-length representation of the utterance response pairs i.e. turn semantics, we use an average pool at the last encoder layer of the BERT token representations. We also tried other approaches such as using the classification token instead of pooling, but based on our initial results simple pooling performed consistently better.

We share our BERT-based LM parameters across the network to encode the session turns. However, we train separate GRU networks to summarize the previous and next turns. The output dimension of the LM is equal to 768, the size of the standard BERT hidden layer. The hidden layer and output size of our GRUs are 256, and we use 2-layer bi-directional GRUs. Each head is a simple MLP with a single hidden layer of size 256 followed by a ReLU nonlinearity. The final network consists of about 117.7 million parameters from which about 110 million is related to BERT and the rest is for GRUs, heads, etc.

### 3.4 Supervised Learning Baseline

As a baseline approach, we use the network defined in Section 3.3 with a binary classification head to distinguish SAT and DSAT samples. Here, we



use labels provided by  $\mathbb{D}_{sup}$  and a binary cross-entropy (BCE) loss function. An Adam optimizer (Kingma and Ba, 2014) with a batch size of 512 is used to train the network for 10 epochs. The base learning rate for all non-BERT layers is set to  $10^{-3}$ , while for BERT layers, we use a smaller learning rate of  $5 \times 10^{-5}$ . The learning rates are decayed with a factor 5 twice at 60% and 80% of total iterations. Unless indicated otherwise, we use a similar training setup for other experiments suggested in this paper.

## 4 Self-Supervised Contrastive Learning

### 4.1 Self-Supervised Objective

We define a self-supervised objective in which the model is tasked to distinguish real sessions from unreal (or noisy) sessions. Any unlabeled dataset, such as  $\mathbb{D}_{un\text{sup}}$  can be used to sample real sessions. To generate unreal textual information, different approaches have been suggested in the literature such as back-translation (Fang and Xie, 2020), generative modeling (Liu et al., 2020), or even random word substitutions.

In this work, we leverage the multi-turn and structured nature of sessions to generate noise samples by simply shuffling the targeted utterances/responses within each training batch (see Figure 3 for an example). Intuitively, the noise samples are sessions in which the targeted utterance or response does not belong to the rest of the session. Therefore, the model has to capture the joint distribution of the context and targeted turns. Algorithm 1 shows an overview of the sample generation and training process for the proposed contrastive objective.

- Sample 1 (+)**    **U:** Play  
                           **R:** What do you want me to play?
- Sample 2 (+)**    **U:** What time is it?  
                           **R:** The time is 12:55 pm
- Sample 3 (-)**    **U:** Play  
                           **R:** The time is 12:55 pm
- Sample 4 (-)**    **U:** What time is it?  
                           **R:** What do you want me to play?

Figure 3: A toy example demonstrating the generation of unreal samples from a batch of two real samples. Session context is omitted for brevity.

---

### Algorithm 1 Contrastive Self-Supervised Training

---

**Input:**  $\mathbb{D}_{un\text{sup}}, h_\theta$  (model w/ contrastive head)  
**repeat**  
    $X \leftarrow \text{GetBatch}(\mathbb{D}_{un\text{sup}})$   
    $batchsize \leftarrow \text{length}(X)$   
    $y \leftarrow \text{ones}(batchsize)$   
    $X_n \leftarrow \text{clone}(X)$   
   **if**  $\text{rand}() < 0.5$   
       $\text{shuffle}(X_n[\text{'targeted\_utterance'}])$   
   **else**  
       $\text{shuffle}(X_n[\text{'targeted\_response'}])$   
    $y_n \leftarrow \text{zeros}(batchsize)$   
    $p \leftarrow h_\theta([X; X_n])$   
    $loss \leftarrow \text{BCE}(p, [y; y_n])$   
   Backprop  $loss$   
   Update  $\theta$   
**until**  $\text{MaxEpoch}$

---

### 4.2 Contrastive Pretraining

The objective introduced in Section 4.1 is not directly applicable to be used as a user satisfaction model. One approach to leverage the pool of unsupervised data is to pre-train the model on unlabeled data using the self-supervised objective, and then attach a classifier head and finetune the network to distinguish SAT and DSAT samples. In our implementation, we pre-train using the self-supervised objective on  $\mathbb{D}_{un\text{sup}}$  for 10 epochs, then train a classifier head on  $\mathbb{D}_{sup}$  for another 10 epochs; adjusting the learning rates for the network body to  $\times 0.1$  of the base learning rates (see Section 3.4 for more information on the learning rate setup).

### 4.3 Few-Shot Learning

In the pretraining approach, we solely relied on the loose semantic relationship between the self-supervised and the user satisfaction modeling tasks. However, it is desirable to have a representation that is not only solving the self-supervised task but is also useful for the final objective. In other words, we have a source task ( $S$ ) which we have a large number of training samples and a target task ( $T$ ) with a limited number of samples that is our main interest. The idea is to use information from the target task during the source training such that the trained model is most compatible with the target.

Let us assume we have datasets  $\mathbb{D}_S$  and  $\mathbb{D}_T$  corresponding to the source ( $S$ ) and target ( $T$ ) tasks as well as inference functions for each task:  $f_S(\cdot|\theta, \omega_S)$  and  $f_T(\cdot|\theta, \omega_T)$ . In this notation,  $\theta$  rep-

resents shared network parameters (i.e., the body in our architecture) and  $\omega$  represents task-specific parameters (i.e., a head in our architecture). Formally, when optimizing for task  $S$ , we are interested in:

$$\arg \min_{\theta, \omega_S} \mathbb{E}_{\mathbf{x}, y \sim \mathbb{D}_S} [L_S(f_S(\mathbf{x}|\theta, \omega_S), y)], \quad (2)$$

where  $L_S$  is the loss function for the source task. A simple gradient descent step to solve this problem can be written as:

$$\begin{aligned} \theta^{t+1} &\leftarrow \theta^t - \eta \nabla_{\theta} \mathbb{E}(L_S(f_S(\mathbf{x}|\theta^t, \omega_S^t), y)), \\ \omega_S^{t+1} &\leftarrow \omega_S^t - \eta \nabla_{\omega_S} \mathbb{E}(L_S(f_S(\mathbf{x}|\theta^t, \omega_S^t), y)). \end{aligned} \quad (3)$$

However, we are interested in optimization steps that do not increase the loss value for task  $T$ :

$$\begin{aligned} \mathbb{E}_{\mathbf{x}, y \sim \mathbb{D}_T} [L_T(f_T(\mathbf{x}|\theta^{t+1}, \omega_T^{t+1}), y)] &\leq \\ \mathbb{E}_{\mathbf{x}, y \sim \mathbb{D}_T} [L_T(f_T(\mathbf{x}|\theta^t, \omega_T^t), y)]. \end{aligned} \quad (4)$$

Considering (4) as an optimization constraint can potentially halt the optimization because improvements to the source objective do not directly translate to improvements to the target task. In other words, the constraint above may not be always directly satisfiable using gradient steps in the source domain.

To overcome this issue, instead of using gradient descent, we define the problem as a Randomize Block Coordinate Descent (RBCD) (Nesterov, 2012; Wright, 2015) optimization. At each RBCD iteration, only a subset of model parameters, i.e. a block noted as  $\mathbf{b}$ , is sampled from a distribution  $\mathcal{B}$  and used for the gradient descent update<sup>2</sup>:

$$\begin{aligned} \mathbf{b} &\sim \mathcal{B}, \\ \theta_{\mathbf{b}}^{t+1} &\leftarrow \theta_{\mathbf{b}}^t - \eta \nabla_{\mathbf{b}} \mathbb{E}(L_S(f_S(\mathbf{x}|\theta^t, \omega_S^t), y)). \end{aligned} \quad (5)$$

Note that we only use the RBCD optimization for the network body parameters ( $\theta$ ), while the head parameters ( $\omega_S$  and  $\omega_T$ ) are optimized using a regular gradient descent optimization.

In this work, we propose the idea of adjusting the block selection distribution,  $\mathcal{B}$ , such that parameters having more aligned source and target gradients have more chance of being selected:

$$\mathcal{B} : Pr(i \in \mathbf{b}) \propto \langle \nabla_{i,S} L_S, \nabla_{i,T} L_T \rangle, \quad (6)$$

where the inputs to  $L_S$  and  $L_T$  are omitted for brevity. Intuitively, (6) is used to discourage parameter updates that are not aligned with the T task

<sup>2</sup>Note that the block selection operation is discrete, either a certain parameter belongs to the block or not, but the distribution  $\mathcal{B}$  can be a continuous or discrete probability distribution.

which can be viewed as a soft method to enforce the constraint in (4). Here, there are multiple options to define the granularity of the block selection such as layer-wise, neuron-wise, or element-wise. Based on our initial experiments, we found that defining the block elements to be layer-wise results in the best performance.

Algorithm 2 shows an outline of the proposed method. At each iteration within the training loop, we back-propagate the  $S$  and  $T$  losses and store the gradients of layer parameters. For parameters related to the  $S$  head, we follow a simple gradient descent update. For body parameters, we only update the parameters if the inner product of the  $S$  and  $T$  tasks is positive or at a small random outcome with the probability of  $\alpha$ . To guarantee the convergence of the source task, we allow all parameters to be selected at each step at least with a very small probability of  $\alpha$ . In our experiments, we consider  $\alpha$  as a hyperparameter taking values in  $\{0.001, 0.005, 0.01, 0.05, 0.1\}$ . Additional care is required when updating the  $T$  head layer parameters as the  $\mathbb{D}_T$  is usually much smaller than  $\mathbb{D}_S$  and the  $T$  head is prone to overfitting. We use a validation set from task  $T$  to detect overfitting for the  $T$  head and early stop the updates. Note that a hyperparameter  $\lambda$  is used to set the frequency of the  $T$  head updates after the early stopping. Having less frequent head updates allows the  $T$  head to gradually improve and adapt to the changes in the body without getting overfitted. In our experiments, we search for proper  $\lambda$  values in  $\{0.001, 0.002, 0.005, 0.01\}$ .

In contrast to other works in the literature which mostly leverage the alignment of concatenated gradients (Lopez-Paz and Ranzato, 2017; Luo et al., 2020), we propose layer-wise similarity measurements providing more granularity and more adaptability. Also, the suggested approach does not require any inner loop optimization process or gradient projection and hence is scalable to large-scale problems. The only computational and memory overhead is to store the model gradients with respect to each task and to compute inner products between the layer parameters.

The method explained in this section is general to few-shot transfer learning and joint training settings where a large source dataset is being used to achieve representations that are most useful for a final target task. For our use-case, we use the suggested approach considering the source task,  $S$ ,

---

**Algorithm 2** The Proposed Few-Shot Training

---

**Input:**  $\mathbb{D}_S, \mathbb{D}_T, f_S, f_T, \alpha$  (random selection rate),  $\lambda$  ( $T$  head update rate)

**repeat**

$(\mathbf{x}_S, y_S) \sim \mathbb{D}_S$

$(\mathbf{x}_T, y_T) \sim \mathbb{D}_T$

// compute & store gradients

$loss_S \leftarrow L_S(f_S(\mathbf{x}_S, y_S))$

Backprop and store  $loss_S$

$loss_T \leftarrow L_T(f_T(\mathbf{x}_T, y_T))$

Backprop and store  $loss_T$

// Layer-Wise RBCD update

**for**  $\mathcal{P}$  in *LayerParameters* :

**if**  $\mathcal{P} \in \omega_S$  // if S head param

$\mathcal{P} \leftarrow \mathcal{P} - \eta \nabla_{\mathcal{P}} loss_S$

**else if**  $\mathcal{P} \in \theta$  // if body param

$sim \leftarrow \langle \nabla_{\mathcal{P}} loss_S, \nabla_{\mathcal{P}} loss_T \rangle$

**if**  $sim > 0$  or  $rand() < \alpha$

$\mathcal{P} \leftarrow \mathcal{P} - \eta \nabla_{\mathcal{P}} loss_S$

**else**

// if T head parameter

**if** *NotEarlyStopped* or  $rand() < \lambda$

$\mathcal{P} \leftarrow \mathcal{P} - \eta \nabla_{\mathcal{P}} loss_T$

Validate, update *NotEarlyStopped*

**until** *MaxEpoch*

---

as the self-supervised contrastive objective and the target task,  $T$ , as the user satisfaction prediction task. In our experiments, after the joint training process, we reinitialize the  $T$  head and finetune the network for the  $T$  task. We found this approach to be helpful to achieve the best results as the jointly trained  $T$  head is often slightly overfitted.

## 5 Results

### 5.1 Experimental Setup

We used PyTorch (Paszke et al., 2017) to train our models. For each case, we continue the training for the maximum number of epochs (10 in our experiments) and select the best model based on the validation performance. We conducted our experiments on a cluster of 48 NVIDIA V100 GPUs (16 GB memory, 8 GPUs per machine). It took between about 6 hours to 27 hours to run individual experiments, depending on the case.

For each experiment, we report Area Under the ROC Curve (AUC-ROC) and Area Under the Precision-Recall Curve (AUC-PR) as the performance measures. The results for the in-domain and out-of-domain held-out test sets are reported

separately. Note that there is an imbalance in the frequency of SAT and DSAT labels, and also there is a difference in the label distribution for the in-domain and out-of-domain test sets. To ensure the statistical significance of the results, each experiment is repeated four times using random initializations reporting the mean and standard deviations.

### 5.2 Quantitative Results

Figure 4 shows a comparison of the in-domain test results for the supervised training and the self-supervised contrastive pretraining methods. For each case, we report the in-domain test performance using models trained with a different number of annotated training samples. The x-axis is plotted in the log scale. It can be seen that the contrastive self-supervised approach is much more data-efficient compared to the supervised approach as it leverages the pool of unlabeled data.

Figure 5 shows a comparison between the supervised training and the self-supervised pretraining methods on the out-of-domain test set. Similar to the in-domain case, there is a significant gap between the labeled data efficiency of these approaches. However, compared to the in-domain case, using even all training samples, the gap does not appear to close. In other words, for the out-of-domain test set the self-supervised approach is not only more data-efficient but also tends to generalize better. In a real-world conversation system, the out-of-domain generalization can be crucial as many different new skills are being developed and included in the system every day, making the traditional in-domain human annotation less practical due to the required annotation turnaround time.

In Figure 6 and Figure 7, we compare the in-domain and out-of-domain performance of the self-supervised pretraining method with the proposed few-shot learning method. As it can be seen from Figure 6, the in-domain AUC-PR and AUC-ROC for the few-shot learning are consistently better than the self-supervised pretraining approach. Note that the performance gap closes at about 5000 samples; perhaps because it is enough training data for fine-tuning and successfully transferring the pre-trained model. The out-of-domain performances as reported in Figure 7 show better results for the few-shot approach but the margin of improvement is relatively smaller than the in-domain case.

Note that in the presented results, we focused our comparisons to methods that are scalable and lever-

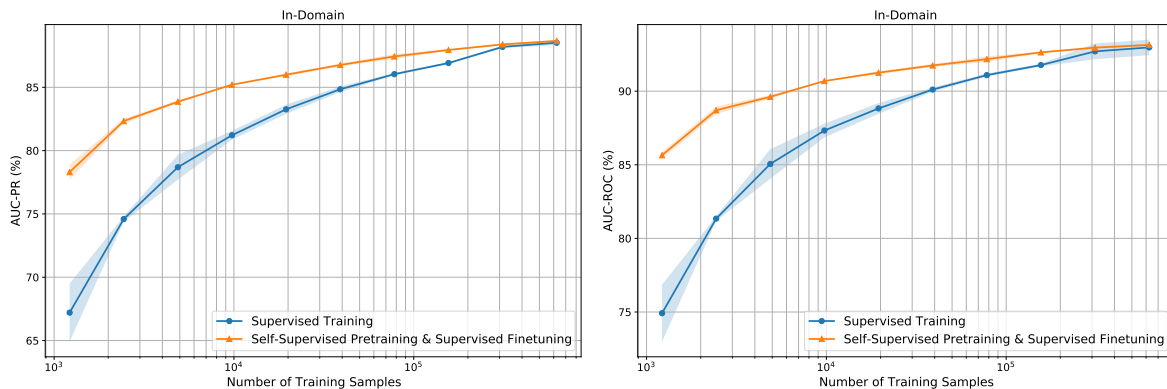


Figure 4: Comparison of the supervised training baseline and the proposed self-supervised pretraining methods for the in-domain test set using different number of training samples (left: AUC-PR, right: AUC-ROC).




Figure 5: Comparison of the supervised training baseline and the proposed self-supervised pretraining methods for the out-of-domain test set using different number of training samples (left: AUC-PR, right: AUC-ROC).

age human annotation data for turn-level satisfaction prediction, excluding approaches using human-engineered and skill-specific metrics as well as methods that only consider the quality of conversation from the language perspective.

### 5.3 Qualitative Results

Table 3 in Appendix B presents a qualitative comparison of the baseline supervised training and the self-supervised approach suggested in this paper. Here, to highlight the generalization and data-efficiency of each method, we limit the number of annotated samples to 1024 random samples from the training set of the  $\mathbb{D}_{sup}$  dataset. For this table, we provide sample sessions that are chosen with an emphasis on more difficult requests, unclear requests, or requests involving 3p skills.  $U$  and  $R$  indicate the targeted utterance and response, while  $U + x$  and  $R + x$  indicate the context utterance and responses appearing  $x$  turns after the targeted turn.

From the provided examples, it can be inferred that the self-supervised approach provides a deeper understanding of the user-agent interaction and is

able to generalize better even for infrequent 3p skills. It is consistent with the quantitative results presented in the paper.

## 6 Conclusion

This paper suggested a self-supervised objective to learn user-agent interactions leveraging large amounts of unlabeled data available. In addition to the standard fine-tuning approach, this paper presented a novel few-shot transfer learning method based on adjusting the RBCD block selection distribution to favor layer parameters with source and target gradients pointing in similar directions. According to the experiments using real-world data, the proposed approach not only requires significantly less number of annotations, but also generalizes better for unseen out-of-domain skills.






Figure 6: Comparison of the self-supervised pretraining and the proposed few-shot learning methods for the in-domain test set using different number of training samples (left: AUC-PR, right: AUC-ROC).




Figure 7: Comparison of the self-supervised pretraining and the proposed few-shot learning methods for the out-of-domain test set using different number of training samples (left: AUC-PR, right: AUC-ROC).

## References

- Praveen Kumar Bodigutla, Lazaros Polymenakos, and Spyros Matsoukas. 2019. Multi-domain conversation quality evaluation via user satisfaction estimation. *arXiv preprint arXiv:1911.08567*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- R Devon et al. 2020. Representation learning with video deep infomax. *arXiv preprint arXiv:2007.13278*.
- Hongchao Fang and Pengtao Xie. 2020. Cert: Contrastive self-supervised learning for language understanding. *arXiv preprint arXiv:2005.12766*.
- Steve Fox, Kuldeep Karnawat, Mark Mydland, Susan Dumais, and Thomas White. 2005. Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems (TOIS)*, 23(2):147–168.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304.
- Ahmed Hassan. 2012. A semi-supervised approach to modeling web search satisfaction. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 275–284.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2018. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*.
- Kalervo Järvelin, Susan L Price, Lois ML Delcambre, and Marianne Lykke Nielsen. 2008. Discounted cumulated gain based evaluation of multiple-query ir sessions. In *European Conference on Information Retrieval*, pages 4–15. Springer.
- Jiepu Jiang, Ahmed Hassan Awadallah, Rosie Jones, Umut Ozertem, Imed Zitouni, Ranjitha Gurunath Kulkarni, and Omar Zia Khan. 2015. Automatic online evaluation of intelligent assistants. In

- Proceedings of the 24th International Conference on World Wide Web*, pages 506–516.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Junnan Li, Ziwei Xu, Wong Yongkang, Qi Zhao, and Mohan Kankanhalli. 2020. Gradmix: Multi-source transfer across domains and tasks. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3019–3027.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Chia-Wei Liu, Ryan Lowe, Iulian V Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. *arXiv preprint arXiv:1603.08023*.
- Xiao Liu, Fanjin Zhang, Zhenyu Hou, Zhaoyu Wang, Li Mian, Jing Zhang, and Jie Tang. 2020. Self-supervised learning: Generative or contrastive. *arXiv*, pages arXiv–2006.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in neural information processing systems*, pages 6467–6476.
- Yan Luo, Yongkang Wong, Mohan S Kankanhalli, and Qi Zhao. 2020.  $n$ -reference transfer learning for saliency prediction. *arXiv preprint arXiv:2007.05104*.
- Yu Nesterov. 2012. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for nlg. *arXiv preprint arXiv:1707.06875*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Dookun Park, Hao Yuan, Dongmin Kim, Yinglei Zhang, Matsoukas Spyros, Young-Bum Kim, Ruhi Sarikaya, Edward Guo, Yuan Ling, Kevin Quinn, et al. 2020. Large-scale hybrid approach for predicting user satisfaction with conversational agents. *arXiv preprint arXiv:2006.07113*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch.
- Louisa Pragst, Stefan Ultes, and Wolfgang Minker. 2017. Recurrent neural network interaction quality estimation. In *Dialogues with Social Robots*, pages 381–393. Springer.
- Niklas Rach, Wolfgang Minker, and Stefan Ultes. 2017. Interaction quality estimation using long short-term memories. In *Proceedings of the 18th Annual SIG-dial Meeting on Discourse and Dialogue*, pages 164–169.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, et al. 2018. Conversational ai: The science behind the alexa prize. *arXiv preprint arXiv:1801.03604*.
- Trieu H Trinh, Minh-Thang Luong, and Quoc V Le. 2019. Selfie: Self-supervised pretraining for image embedding. *arXiv preprint arXiv:1906.02940*.
- Stephen J Wright. 2015. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34.
- Tiansheng Yao, Xinyang Yi, Derek Zhiyuan Cheng, Felix Yu, Aditya Menon, Lichan Hong, Ed H Chi, Steve Tjoa, Evan Ettinger, et al. 2020. Self-supervised learning for deep models in recommendations. *arXiv preprint arXiv:2007.12865*.

## A Annotation Protocol

In the following, we provide a summary of main points considered to produce the annotations used in this paper<sup>3</sup>:

- Human annotators were trained to annotate samples i.e., we do not use domain specific metrics or other automated success measures as annotation.
- It was made clear to the annotators that the task is turn-level user satisfaction, and not the overall satisfaction over session. Also, instructions were provided on how to handle ASR errors, repeated requests, multiple users in one utterance, and many other special cases.
- The annotators were provided the targeted turn as well as a few context turns. This helped them to better understand the actual user intention and judge accordingly.
- They were asked to rate the system’s response quality in terms of the user satisfaction on the scale of 1 to 5, from terrible to excellent. See Table 2 for score categories and an example of each category.
- To ensure the quality of annotations, each sample was annotated multiple times by different annotators.

Table 2: Examples of the annotation scores.

Score	Example
1 (Terrible)	U: <i>Is bacon really bad for your health?</i> R: <i>Sorry, I'm not sure.</i>
2 (Bad)	U: <i>How's the humidity?</i> R: <i>Sorry, I don't have daily forecast information.</i>
3 (Okay)	U: <i>Play a song for the baby.</i> R: <i>Skill [SKILL NAME] can help you with that.</i> <i>Did you want to enable it?</i>
4 (Good)	U: <i>What's the temperature going to be tomorrow?</i> R: <i>Tomorrow in London, you can look for cloudy weather with a high of 65 and low of 52 degrees.</i>
5 (Excellent)	U: <i>Cancel alarm.</i> R: <i>6:35 PM alarm cancelled.</i>

In our analysis, we consider all samples having a score of 3 or better as SAT, and DSAT otherwise. Also, in our data pipeline, we considered different annotations of the same utterance as different samples. However, care was taken in the data split

<sup>3</sup>Certain details were omitted to comply with confidentiality requirements.

process to ensure there is no train data contamination in our validation and test sets.

## B Qualitative Results

Table 3 presents a qualitative comparison of the baseline supervised training and the self-supervised approach suggested in this paper. Here, to highlight the generalization and data-efficiency of each method, we limit the number of annotated samples to 1024 random samples from the training set of the  $\mathbb{D}_{sup}$  dataset. For this table, we provide sample sessions that are chosen with an emphasis on more difficult requests, unclear requests, or requests involving 3p skills.  $U$  and  $R$  indicate the targeted utterance and response, while  $U+x$  and  $R+x$  indicate the context utterance and responses appearing  $x$  turns after the targeted turn.

From these examples, it can be inferred that the self-supervised approach provides a deeper understanding of the user-agent interaction and is able to generalize better even for infrequent 3p skills. It is consistent with the quantitative results presented in the paper.

Table 3: Qualitative comparison of predictions for different methods on a set of selected test samples.

Session	Label	Prediction		Comments
		Sup <sup>a</sup>	Self <sup>b</sup>	
<p>U-1: <i>what what high school basketball team do you like</i></p> <p>U : <i>what is your favorite m. l. b. team</i></p> <p>U+1: <i>what's your favorite movie</i></p>		SAT	DSAT SAT	Asking for favorite, agent using advanced language and slang
<p>U-1: <i>is it going to rain tomorrow</i></p> <p>U : <i>is it raining now in new york</i></p> <p>U+1: <i>is it raining today</i></p> <p>U+2: <i>set an alarm for five p.m.</i></p>		DSAT	SAT DSAT	Asking for weather, rewording and repeating the request
<p>U : <i>close the garage</i></p>		SAT	DSAT SAT	Device interaction, using a less common 3p skill
<p>U : <i>teach me</i></p> <p>U+1: <i>pause</i></p>		DSAT	SAT DSAT	Knowledge question, vague request agent fails to understand
<p>U-2: <i>what is in the middle of the milky way galaxy</i></p> <p>U-1: <i>name all the countries in the world</i></p> <p>U : <i>read us a really long list</i></p> <p>U+1: <i>cancel</i></p> <p>U+2: <i>name all the countries in europe</i></p>		DSAT	SAT DSAT	Knowledge question, agent fails to understand user's intention
<p>U-2: <i>call</i></p> <p>U-1: <i>p.m</i></p> <p>U : <i>call liam</i></p> <p>U+1: <i>thank you [AGENT]</i></p>		SAT	DSAT SAT	Call request, failed due to system limitation but user understands

<sup>a</sup>Supervised training baseline

<sup>b</sup>Self-supervised contrastive learning method