

Conversation Model Fine-Tuning for Classifying Client Utterances in Counseling Dialogues

Sungjoon Park¹, Donghyun Kim², Alice Oh¹

¹ School of Computing, KAIST, Republic of Korea

² Trost, Humart Company, Inc., Republic of Korea

sungjoon.park@kaist.ac.kr, kdhyeon821@gmail.com

alice.oh@kaist.edu

Abstract

The recent surge of text-based online counseling applications enables us to collect and analyze interactions between counselors and clients. A dataset of those interactions can be used to learn to automatically classify the client utterances into categories that help counselors in diagnosing client status and predicting counseling outcome. With proper anonymization, we collect counselor-client dialogues, define meaningful categories of client utterances with professional counselors, and develop a novel neural network model for classifying the client utterances. The central idea of our model, ConvMFiT, is a pre-trained conversation model which consists of a general language model built from an out-of-domain corpus and two role-specific language models built from unlabeled in-domain dialogues. The classification result shows that ConvMFiT outperforms state-of-the-art comparison models. Further, the attention weights in the learned model confirm that the model finds expected linguistic patterns for each category.

1 Introduction

Some mental disorders are known to be treated effectively through psychotherapy. However, people in need of psychotherapy may find it challenging to visit traditional counseling services because of time, money, emotional barriers, and social stigma (Bearse et al., 2013). Recently, technology-mediated psychotherapy services emerged to alleviate these barriers. Mobile-based psychotherapy programs (Mantani et al., 2017), fully automated chatbots (Ly et al., 2017; Fitzpatrick et al., 2017), and intervention through smart devices (Torrado et al., 2017) are examples. Among them, text-based online counseling services with professional counselors are becoming popular because clients can receive these services without traveling to an

office and with reduced financial burden compared to traditional face-to-face counseling sessions (Hull, 2015).

In text-based counseling, the communication environment changes from face-to-face counseling sessions. The counselor cannot read non-verbal cues from their clients, and the client uses text messages rather than spoken utterances to deliver their thoughts and feelings, resulting in changes of dynamics in the counseling relationship. Previous studies explored computational approaches to analyzing the dynamic patterns of relationship between the counselor and the client by focusing on the language of counselors (Imel et al., 2015; Althoff et al., 2016), clustering topics of client issues (Dinakar et al., 2014), and looking at therapy outcomes (Howes et al., 2014; Hull, 2015).

Unlike previous studies, we take a computational approach to analyze client responses from the *counselor's perspective*. Client responses in counseling are crucial factors for judging the counseling outcome and for understanding the status of the client. So we build a novel categorization scheme of client utterances, and we base our categorization scheme on the cognitive behavioral theory (CBT), a widely used theory in psychotherapy. Also, in developing the categories, we consider whether they are adequate for the unique text-only communication environment, and appropriate for the annotation of the dialogues as training data. Then using the corpus of text-based counseling sessions annotated according to the categorization scheme, we build a novel conversation model to classify the client utterances.

This paper presents the following contributions:

- First, we build a novel categorization method as a labeling scheme for client utterances in text-based counseling dialogues.

- Second, we propose a new model, Conversation Model Fine-Tuning (ConvMFiT) to classify the utterances. We explicitly integrate pre-trained language-specific word embeddings, language models and a conversation model to take advantage of the pre-trained knowledge in our model.
- Third, we empirically evaluate our model in comparison with other models including a state-of-the-art neural network text classification model. Also, we show typical phrases of counselors and clients for each category by investigating the attention layers.

2 Categorization of Client Utterances

Client responses provide essential clues to understanding the client’s internal status which can vary throughout counseling sessions. For example, client’s responses describing their problems prevail at the early stage of counseling (Hill, 1978; E. Hill et al., 1983), but as counseling progresses, problem descriptions decrease while insights and discussions of plans continue to increase (Seeman, 1949). Client responses can also help in predicting counseling outcomes. For example, a higher proportion of insights and plans in client utterances indicates a high positive effect of counseling (Hill, 1978; E. Hill et al., 1983).

Categorization Objective. Our final aim is to build a machine learning model to classify the client utterances. Thus, the categorization of the utterances should satisfy the following criteria:

- *Suitable for the text-only environment:* Categories should be detected only using the text response of clients.
- *Available as a labeling scheme:* The number of categories should be small enough for manual annotation by counseling experts.
- *Meaningful to counselors:* Categories should be meaningful for outcome prediction or counseling progress tracking.

Previous studies in psychology proposed nine and fourteen categories for client and counselor verbal responses, respectively, by analyzing transcriptions from traditional face-to-face counseling sessions (Hill, 1978; Hill et al., 1981). But these categories were developed for face-to-face spoken interactions, and we found that for online text-only

counseling dialogues, these categories are not directly applicable. Using text without non-verbal cues, a client’s responses are inherently different from the transcriptions of verbally spoken responses which include categories such as ‘silence’ (no response for more than 5 seconds) and ‘non-verbal referent’ (physically pointing at a person). Another relevant study, derived from text-based counseling sessions with suicidal adolescents proposes 19 categories which we judged to be too many to be practical for manual annotation (Kim, 2010).

The last criterion of “meaningful to counselors” is perhaps the most important. To meet that criterion, we base the categorization process on the *cognitive behavioral theory (CBT)* which is the underlying theory behind psychotherapy counseling. The details of using CBT for the categorization process is explained next.

Categorization Process and Results. In developing the categories, we follow the Consensual Qualitative Research method (Hill et al., 1997). Two professional counselors with clinical experience participated in this qualitative research method to define the categorization.

To begin, we randomly sample ten client cases considering demographic information including age, gender, education, job, and previous counseling experiences. We then start the categorization process with the fundamental components of the CBT which are *events, thoughts, emotions, and behavior* (Hill et al., 1981). The professional counselors annotate every client utterance to with those initial component categories with tags that add detail. For example, if an utterance is annotated as ‘emotion’, we add ‘positive/negative’ or concrete label such as ‘hope’. If these tags are categorized to be a new category, we add that category to the list until it is saturated. When the number of categories becomes more than 40, we define higher level categories that cover the existing categories.

In the second stage, annotators discuss and merge these categories into five high-level categories. Category 1 is *informative responses to counselors*, and category 2 is *providing factual information and experiences*. Categories 3 and 4 are related to the client factors, *expressing appealing problems* and *psychological changes*. The last category is about the *logistics of the counseling sessions* including scheduling the next session. The categories in detail are as follows:

Characteristic	Informative		Client Factors		Process
Category Name	Factual Information (Fact.)	Anecdotal Experience (Anec.)	Appealing Problem (Prob.)	Psychological Change (Chan.)	Counseling Process (Proc.)
Explanation	Brief mention of categorical information	Clients experience contributing to the appealing problem	Clients factors related to the appealing problem	Statement at the resolution stage of the appealing problem	Statement of counseling structure and relationship
Examples	<ul style="list-style-type: none"> • Objective Fact • Living conditions • Demographic information • Limited conditions 	<ul style="list-style-type: none"> • Experience with others • Comments from others • Trauma • Interpersonal situations 	<ul style="list-style-type: none"> • Negative Emotion • Cognitive distortion • Interpersonal problems • Family problems 	<ul style="list-style-type: none"> • Positive Prediction • Expectation, Determination • Coping behaviors • Self-awareness 	<ul style="list-style-type: none"> • A message to counselor • Gratitude, Greetings • Time appointment • Questions about the consultation

Table 1: Final Categorization of client utterances. Five categories are discovered, two for informative giving information to a counselor (Factual information, Anecdotal Experience), two for client factors (appealing problems, psychological change), and the last one for counseling process.

- *Factual information (Fact.)* Informative responses to counselor’s utterances, including age, gender, occupation, education, family, previous counseling experience, etc.
- *Anecdotal Experience (Anec.)* Responses describing past incidents and current situations related to the formation of appealing problems. Responses include traumatic experiences, interactions with other people, comments from other people, and other anecdotal experiences.
- *Appealing Problems (Prob.)* Utterances addressing the main appealing problem which is yet to be resolved, including client’s internal factors or their behaviors related to the problems. Specifically, the utterances include cognition, emotion, physiological reaction, and diagnostic features of the problem and desire to be changed.
- *Psychological Change (Chan.)* Utterances describing insights, cognition of small and big changes in internal factors or behaviors. That is, an utterance at the point where the appealing problem is being resolved.
- *Counseling Processes (Proc.)* Utterances that include the objective of counseling, requests to the counselor, plans about the counseling sessions, and counseling relationship. This

category also covers greetings and making an appointment for the next session.

We summarize the category explanations and examples in Table 1.

3 Dataset

In this section, we explain how counseling dialogues differ from general dialogues, describe the dialogues we collected and annotated, and explain how we preprocessed the data.

3.1 Characteristics of Counseling Dialogues

The counseling dialogues consist of multiple turns taken by a counselor and a client, and each turn can contain multiple utterances. Here we describe two unique characteristics of text-based online counseling conversations compared to general non-goal oriented conversations.

Distinctive roles of speakers. Counseling conversations are goal-oriented with the aim to produce positive counseling outcomes, and the two speakers have distinctive roles. The client gives objective information about themselves and subjective experiences and feelings to the counselor to appeal the problems they are suffering from. Then the counselor establishes a therapeutic relationship with the client and elicits various strategies to induce psychological changes in the client. These distinct roles of the conversational participants distinguish counseling conversations from

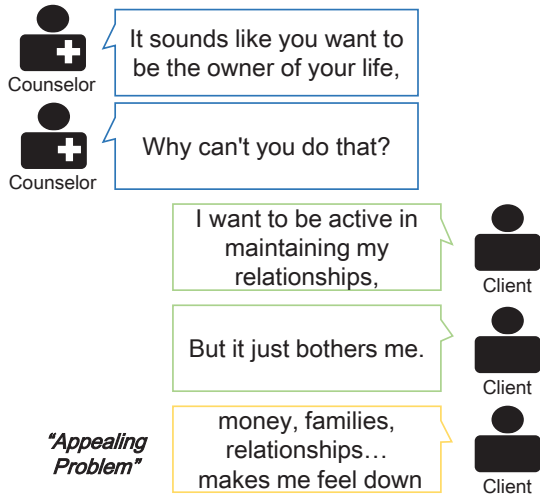


Figure 1: Translated example conversation between a counselor and a client. Each turn can have multiple utterances, and we annotate the client turn at the *utterance*-level. The first two boxes indicate 1) *Counselor’s utterance* and the next two are 2) *Client’s context utterance*, and the last box is 3) *Client’s target utterance*. The annotation for the last client utterance is ”appealing problem”.

non-goal oriented open-domain response generation modeling.

Multiple utterances in a turn. We define an *utterance* as a single text bubble. Counselors and clients can generate multiple utterances in a turn. Especially in a client turn, various information not to be missed by a counselor may occur across multiple utterances. Thus, we treat every utterance separately, as shown in Fig. 1

3.2 Collected Dataset

Total dialogues. We collect counseling dialogues of clients with their corresponding professional counselors from Korean text-based online counseling platform Trost.¹ Overall, we use 1,448 counseling dialogues which are anonymized before researchers obtain access to the data, by removing personally identifiable information in the dialogues. All meta-data of the dialogues are not provided to the researchers, and named entities such as client’s and counselor’s names are replaced with random numeric identifiers. The research process including data anonymization and pre-processing is validated by the KAIST Institutional Review Board (IRB).

Labeled Dialogues. We randomly choose and label 100 dialogues with the discovered five categories. Note that we only label client utterances.

¹<https://www.trost.co.kr/>

Based on these categories, five professional counselors annotated their own client’s every utterance in the conversations, as shown in Fig. 1. Note that each utterance can have multiple labels if it includes information across multiple categories.

Table. 2 shows the descriptive statistics of our labeled dataset. The first two rows present the average lengths of each utterance of counselors and clients in terms of words and characters, showing there is a small difference between the counselor utterance and the client utterance. On the other hand, the average number of utterances in a single counseling session differs; on average, clients write more utterances than counselors.

	Counselor		Client	
	Mean	Std.	Mean	Std.
# of words	6.26	6.25	5.91	14.62
# of chars	25.71	23.44	24.31	61.10
# of utters	163.2	236.97	238.72	578.49

Table 2: Descriptive statistics of labeled counseling dialogues.

3.3 Preprocessing

We intentionally leave in punctuations and emojis since they can help to infer the categories of the client utterances, treating them as separate tokens. Then we construct triples from the labeled dialogues consisting of 1) counselor’s utterances (blue in Fig. 1), 2) client’s context utterances (green), and 3) client’s target utterances to be categorized. (yellow)

We split the dataset into train, validation, and test sets. Table. 3 shows the number of triples in each set, showing factual information (Fact.) and psychological change (Chan.) categories appear less frequently than the others.

4 Model

We introduce ConvMFiT (Conversation Model Fine-Tuning), fine-tuning pre-trained seq2seq based conversation model to classify the client’s utterances.

Background. Our corpus of 100 labeled conversations is not large enough to fully capture the linguistic patterns of the categories without external knowledge. The small size of the dataset is difficult to overcome because the labeling by professional counselors is costly. We found a poten-

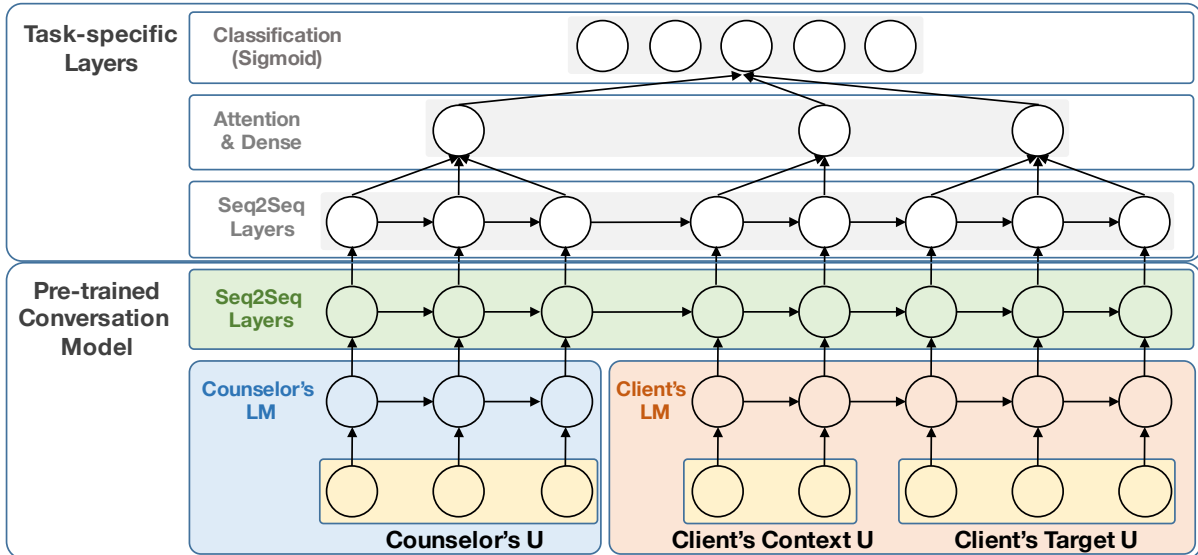


Figure 2: ConvMFiT (Conversation Model Fine-Tuning) model architecture. We use a pre-trained conversation model using seq2seq models. The model is based on pre-trained counselor’s and client’s language models. (lower colored part) Then, we stack additional task-specific seq2seq layers and classification layers over the conversation model to learn specific features. Between the two layers, attention layer is added to enhance its interpretability.

	Train	Valid	Test	# of labels
Fact.	817	140	172	1129
Anec.	5317	1180	1175	7726
Prob.	4728	1004	981	6713
Chan.	887	211	199	1297
Pros.	4570	964	961	6459
# of triples	14679	3166	3165	21100

Table 3: The number of triples (partner’s utterance, client’s context utterance, client’s target utterance) and corresponding labels for each set. Factual information (Fact.) and psychological change (Chan.) categories have less number of instances compared to the others.

tially effective solution for a small labeled dataset by using pre-trained language models for various NLP tasks (Ramachandran et al., 2017; Howard and Ruder, 2018). Therefore, we focus on transferring knowledge from unlabeled in-domain dialogues as well as a general out-of-domain corpus.

Overview. We illustrate the overall model architecture. We first stack pre-trained seq2seq layers which represent a conversation model (see Fig. 2, lower colored part). Then we stack additional seq2seq layers and classification layers over it to capture task-specific features, and an attention layer is added between the two to enhance the model’s interpretability (see Fig. 2, upper white part).

This approach leverages the advantages of us-

ing a language model based conversation model for transfer learning. The model can be regarded as an extended version of ULMFiT (Howard and Ruder, 2018) with modifications to fit our task. In ConVMFiT, the model accepts a pre-trained seq2seq conversation model that requires two pre-trained language models, like an encoder and a decoder, learning the dependencies between them on the seq2seq layers. This approach is shown to be effective for machine translation, applying a source language model for the encoder and target language model for the decoder (Ramachandran et al., 2017).

4.1 Model Components

Word Vectors. Counseling dialogues consist of natural Korean text which is morphologically rich, so we train word vectors specifically developed for the Korean language (Park et al., 2018a). We train the vectors over a Korean corpus including out-of-domain general documents, 1) Korean Wikipedia, 2) online news articles, and 3) Sejong Corpus, as well as in-domain unlabeled counseling dialogues. The corpus contains 0.13 billion tokens. To validate the trained vector quality, we check performance on word similarity task (WS353) for Korean (Park et al., 2018a). The Spearman correlation is 0.682, which is comparable to the state-of-the-art performance. These vectors are used as inputs (see Fig. 2, yellow).

Pre-trained Language Models. We assume that counselors and clients have different language models because they have distinctive roles in the dialogue. Therefore, we train a counselor language model and a client language model separately. We collect counselor utterances in the total dialogue dataset, except dialogues in the test set, to train a counselor language model, and all others are used for training a client language model. Then, We fine-tune the two trained LMs with utterances in the labeled dialogues.

For each model, we train word-level language models by using multilayer LSTMs. We apply various regularization techniques, weight tying (Inan et al., 2017), embedding dropout and variational dropout (Gal and Ghahramani, 2016), which are used to regularize LSTM language models (Inan et al., 2017; Ramachandran et al., 2017; Merity et al., 2018).

We use a 3-layer LSTM model having 300 hidden units for every layer. We set embedding dropout and output dropout for each layer to 0.2, 0.1, respectively. The pre-trained language models generate inputs to seq2seq layer of a conversation model (Fig. 2, blue and red).

Pre-trained Conversation Model. Next, we train a seq2seq conversation model (Vinyals and Le, 2015). We use the pre-trained counselor language model as an encoder, and the client language model as a decoder. The dependency of the decoder on the encoder is trained by seq2seq layers, stacked over the pre-trained models. (Fig. 2, green)

We stack 2-layer LSTMs over the pre-trained counselor and client language models, respectively. The final states of the LSTMs on the counselor language model is used as an initial state of the LSTMs on the client language model. We set the hidden unit size of 300 for every LSTMs and set output dropout to 0.05. The outputs of pre-trained conversation models is used for inputs to seq2seq layers of the task-specific layers.

During training the conversation model, we regularize the parameters of the model by adding cross entropy losses of the pre-trained counselor and client language models to seq2seq cross entropy loss of the conversation model. Three losses are weighted equally. This prevents catastrophic forgetting of the pre-trained language models and is important to achieve high performance (Ramachandran et al., 2017). Also, there is room

for improvement of the architecture of a conversation model, which integrates pre-trained language models, to capture dialogue patterns better and thus leads to higher classification performance. We will explore the architecture in future work.

Task-specific layers. By leveraging pre-trained language model based conversation model, we finally add layers for classification. In order to capture task-specific features, we first stack seq2seq layers over the conversation model. Then we add attention mechanism for document classification (Yang et al., 2016). Lastly, we use a sigmoid function as an output layer to predict whether the information is included in utterances because multiple categories can appear in a single utterance. (Fig. 2, gray)

We use a 2-layer LSTM model for the seq2seq layers. We set 300 hidden units for every LSTMs and set output dropout to 0.05. The size of attention layer is set to 500.

4.2 Model Training

Thus the model is trained by three steps: 1) training word vectors and two language models, 2) training seq2seq conversation model with pre-trained LMs, and 3) fine-tuning task-specific classification layers, after removing softmax of the conversation model. In the last step, we compute binary logistic loss between predicted probability for each category and label as a loss function. We use Adam with default parameters for the optimizer, in order to train language models, conversation model, and fine-tuning the classifier. Also, gradual unfreezing is applied while training the model, starting updating parameters from the task-specific layers and unfreezing the next lower frozen layer for each epoch. We unfreeze layers every other epoch until all layers are tuned, and we stop training when the validation loss is minimized. All hyperparameters are tuned over the development set.

5 Experiments

5.1 Comparison Models

We compare our model with baseline models in Table 4. Models 1-5 are classifiers which only use the target client utterance to classify it, and models 6-8 are conversation model-based classifiers considering counselor’s utterances and client’s context utterances. All models use the same pre-trained word vectors for a fair comparison.

(1) Random Forest, (2) SVM with RBF kernel. We represent an utterance by an average of all word vectors in it, then feed it to the classifier input.

(3) CNN for text classification (Kim, 2014). In the convolution layer, we set the filter size to 1-10, and 30 filters each, then we apply max-over-time pooling. Then, a dense layer and sigmoid activation is applied over the layer.

(4) RNN Bidirectional LSTM is used where the final states for each direction are concatenated to represent the utterances. Then, a dense layer and a sigmoid activation are stacked.

(5) ULMFiT. A pre-trained client language model is used as a universal language model. Details are the same as described in Section 5.3. In addition, 2-LSTM layers and dense layer with sigmoid activations are stacked over the LM for classification. Gradual unfreezing is applied during training (Howard and Ruder, 2018).

(6) Seq2Seq. Like encoder-decoder based conversation model (Vinyals and Le, 2015), three LSTMs are assigned for each Counselor’s utterances, client’s context/target utterances. The initial states of the client language models are set to the final states of preceding utterances. Then dense & sigmoid layers for classification are stacked over the final state of the client’s target utterances.

(7) HRED. Hierarchical encoder-decoder model (HRED) is used as a conversation model (Serban et al., 2016). For the encoder RNN, counselor’s utterances and client’s context utterances are given as inputs, and their information is stored in context RNN, which delivers it to the decoder accepting client’s target utterances. Like (6) Seq2Seq, dense & sigmoid layers for classification are stacked over the final state of the client’s target utterances.

5.2 Ablation Study

We conduct an ablation study to investigate the effect of the pre-trained models in Table 5.

(1-4) Adding pre-trained models. Model 1-4 have the same architecture as ConvMFiT, which is Model (8) in Table. 4. Model (1) in Table. 5 initializes every parameter randomly. Model (2) starts training only with pre-trained word vectors. Model (3) leverages counselor and client language models as well, and Model (4) shows the performance of ConvMFiT. As Model (3) and (4) use more than two pre-trained components, gradual

unfreezing is applied, unfreezing shallower layers first during training.

(4-1) Task-specific Seq2Seq Layers. Model (4-1) removes task-specific Seq2Seq layers in the model, which leaves only attention and dense layer to capture task-specific features. It may result in an insufficient model capacity to capture relevant features for the task.

(4-2) Effect of Gradual Unfreezing. Model (4-2) is trained without gradual unfreezing, allowing the parameters of every layer in the model change by the gradients from the first epoch.

6 Results

6.1 Classification Performance

We show the performance of our model and comparison models in Table. 4. (1) Random Forests and (2) Support Vector Machines underperform to classify utterances correctly which belong to rarely occurred classes. Compared to (1) and (2), (3) CNN and (4) RNN show better performance since they look at the sequence of words (.416, .431, respectively). RNN shows slightly better performance than CNN. (6) ULMFiT outperforms the others by using a pre-trained client language models (.455).

The client target utterances have their context, and they also depend on the counselor’s preceding utterances, so using the preceding counselor utterance as well as the client context utterances helps to improve the classification performance. When integrating the information using simple (6) seq2seq model, it shows better performance (.530) than (5) ULMFiT. (7) HRED adds a higher-level RNN to seq2seq models, but we find there is little performance gain (.001) which makes the model overfit easily.

(8) ConvMFiT employs pre-trained conversation model based on pre-trained LMs and so outperforms all other baseline models (.642). This is because ConvMFiT integrates conversational contexts and the counselor language model, which helps to capture the patterns of client’s language better. Also, the improvement is higher for the class (Fact.) and (Chan.) which have small numbers of examples since the ConvMFiT could leverage pre-trained knowledge to classify them.

6.2 Effect of Pre-trained Components

With an ablation study applying pre-trained components step by step to our model, we show the

No.	Dep.	Model	F1 (Fact.)	F1 (Anec.)	F1 (Prob.)	F1 (Chan.)	F1 (Proc.)	Macro Prec	Macro Rec	Macro F1
1	X	RF	.000	.564	.420	.000	.723	.476	.269	.341
2	X	SVM(rbf)	.012	.683	.457	.000	.766	.602	.385	.384
3	X	CNN	.211	.528	.506	.128	.706	.450	.397	.416
4	X	RNN	.193	.574	.570	.046	.770	.607	.375	.431
5	X	ULMFiT	.205	.641	.591	.057	.784	.613	.413	.455
6	O	Seq2Seq	.263	.662	.678	.226	.823	.695	.472	.530
7	O	HRED	.261	.706	.675	.193	.820	.680	.475	.531
8	O	ConvMFiT	.441	.761	.726	.447	.835	.716	.602	.642

Table 4: Classification results. Among the models 1-6 which only use client’s utterances to predict categories, (6) ULMFiT show better performance to the others. Models 7-9 are classifiers based on the conversation model, and ConvMFiT outperforms the others (.642).

No.	Emb.	LM	Conv. seq2seq	Grad. Unf.	Task. seq2seq	F1 (Fact.)	F1 (Anec.)	F1 (Prob.)	F1 (Chan.)	F1 (Proc.)	Macro Prec	Macro Rec	Macro F1
1	X	X	X	-	O	.032	.706	.602	.222	.711	.418	.575	.455
2	O	X	X	-	O	.043	.758	.661	.258	.748	.459	.662	.494
3	O	O	X	O	O	.425	.782	.727	.365	.831	.587	.719	.626
4	O	O	O	O	O	.441	.761	.726	.447	.835	.602	.716	.642
4-1	O	O	O	O	X	.307	.738	.666	.304	.802	.513	.687	.563
4-2	O	O	O	X	O	.417	.768	.721	.399	.824	.591	.695	.626
4	O	O	O	O	O	.441	.761	.726	.447	.835	.602	.716	.642

Table 5: Ablation study results. All models use the same architecture of ConvMFiT. Adding pre-trained word vectors (2), counselor and client language models (3), seq2seq layers of conversation model (4) results in a performance improvement. Also, adding task-specific seq2seq layers to ensure the model’s sufficient capacity for capturing relevant features (4-1), and applying gradual unfreezing lead to performance improvement (4-2).

sources of performance improvement in Table 5. Model (1) uses the same architecture but no pre-trained models are applied, showing poor performance due to overfitting (.455). The f1 score is lower than that of the simple seq2seq model. Model (2) only uses pre-trained word vectors, and it helps to increase the performance slightly (.494). Also, Model (3) adds two pre-trained LMs with gradual unfreezing, resulting in a substantial performance increase (.626.), so we find pre-trained LMs are essential to the improvement. Lastly, Model (4) adds the dependency between two language models to fully leverage the pre-trained conversation model. This also helps classify utterances better (.642).

In addition, we find only adding attention and dense layers are not sufficient to learn task-specific features, so providing the model more capacity helps improving the performance. Without task-specific seq2seq layers, model (4-1) shows decreased f1 score (.563). Meanwhile, model (4-2) shows careful training scheme affects to the per-

formance as well.

6.3 Qualitative Results

To investigate how linguistic patterns of counselors and clients differ in the various categories, we report qualitative results based on the activation values of the attention layer. To protect the anonymity of the clients, we explore key phrases from utterances rather than publish parts of the conversation in any form.

To this end, we compute the relative importance of n-grams. For any n of n-gram in an utterance of length N where $n < N$, the relative importance r is computed as follows:

$$\left(\prod_{i=1}^n a_i - (1/N)^n \right) / (1/N)^n \quad (1)$$

where a_i is corresponding attention value of a word, $\prod_{i=1}^n a_i$ is the product of the values for every words in the n-gram, normalized by the expected attention weights $(1/N)^n$. The normalization term considers the length of utterance since a word in short utterances tends to have high atten-

tion value because of $\sum_{i=1}^N a_i = 1$.

We name this measure ‘relative importance’ r meaning that the degree of the n-gram is how much it is attended to compared to the expectation. Based on this, we select examples from 100 top-ranked key phrases for each category and the results are presented in Appendix (Table. 6). All of the presented key phrases are translated to English from Korean.

Factual Information. Clients provide demographic information and previous experience of visits to counselors or psychiatrists. In some case, clients talk more about the motivation of counseling. Since this information is explored in an early stage of counseling sessions, we also find counselor greetings with client’s names.

Anecdotal Experience. Clients describe their experiences by using past tense verbs. Usually, utterances include phrases such as ‘I thought that’, ‘I was totally wrong’. Counselors show simple responses ‘well..’, and reflections.

Appealing Problem. Like anecdotal experience, clients describe their problems, but with using the present tense of verbs. They are appealing their thinking and emotions. Counselors also show simple responses or reflections. Since some clients immediately start pouring out their problems right after counseling sessions starts, so greetings from counselor appear in key phrases.

Psychological Change. Clients obviously report their change of feelings, emotions or thoughts. It includes looking back on the past and then determining to change in the future. Counselors give supportive responses and empathetic understanding.

Counseling Process. Clients and counselors exchange greetings with each other. Also, they discuss making an appointment for the next session. In some cases, counselors respond to client’s questions about the logistics of the sessions.

7 Related Work

Researchers have explored psycho-linguistic patterns of people with mental health problems (Gkotsis et al., 2016), depression (Resnik et al., 2015), Asperger’s and autism (Ji et al., 2014) and Alzheimer’s disease (Orimaye et al., 2014). In addition, these linguistic patterns can be quantified, for example, overall mental health (Loveys et al., 2017; Coppersmith et al., 2014), and schizophrenia (Mitchell et al., 2015).

To aid people with those mental issues, large portion of studies are dedicated to detecting those issues from natural language. Depression (Morales et al., 2017; Jamil et al., 2017; Fraser et al., 2016), anxiety (Shen and Rudzicz, 2017), distress (Desmet et al., 2016), and self-harm risk (Yates et al., 2017) can be effectively detected from narratives or social media postings.

8 Discussion and Conclusion

In this paper, we developed five categories of client utterances and built a labeled corpus of counseling dialogue. Then we developed the ConvMFiT for classifying the client utterances into the five categories, leveraging a pre-trained conversation model. Our model outperformed comparison models, and this is because of transferring knowledge from the pre-trained models. We also explored and showed typical linguistic patterns of counselors and clients for each category.

Our ConvMFiT model will be useful in other classification tasks based on dialogues. ConvMFiT is a seq2seq model for counselor-client conversation, however, another approach would be to model with existing non-goal oriented conversation models incorporating Variational Autoencoder (VAE) (Serban et al., 2017; Park et al., 2018b; Du et al., 2018). We plan to attempt these models in future work.

We expect to apply our trained model to various text-based psychotherapy applications, such as extracting and summarizing counseling dialogues or using the information to build a model addressing the privacy issue of training data. We hope our categorization scheme and our ConvMFiT model become a stepping stone for future computational psychotherapy research.

Ethical Approval

The study reported in this paper was approved by the KAIST Institutional Review Board (#IRB-17-95).

Acknowledgments

This research was supported by the Engineering Research Center Program through the National Research Foundation of Korea (NRF) funded by the Korean Government MSIT (NRF-2018R1A5A1059921).

References

- Tim Althoff, Kevin Clark, and Jure Leskovec. 2016. [Large-scale analysis of counseling conversations: An application of natural language processing to mental health](#). *Transactions of the ACL*, 4:463–476.
- Jennifer L Bearse, Mark R McMinn, Winston See-gobin, and Kurt Free. 2013. Barriers to psychologists seeking mental health care. *Professional Psychology: Research and Practice*, 44(3):150.
- Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. [Quantifying mental health signals in twitter](#). In *Proceedings of the Workshop on CLPsych*. ACL.
- Bart Desmet, Gilles Jacobs, and Véronique Hoste. 2016. [Mental distress detection and triage in forum posts: The lt3 clpsych 2016 shared task system](#). In *Proceedings of the Third Workshop on CLPsych*. ACL.
- Karthik Dinakar, Allison JB Chaney, Henry Lieberman, and David M Blei. 2014. Real-time topic models for crisis counseling. In *Twentieth ACM Conference on Knowledge Discovery and Data Mining, Data Science for the Social Good Workshop*.
- Jiachen Du, Wenjie Li, Yulan He, Ruifeng Xu, Lidong Bing, and Xuan Wang. 2018. Variational autoregressive decoder for neural response generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3154–3163.
- Clara E. Hill, Jean A. Carter, and Mary K. O’Farrell. 1983. Reply to howard and lambert: Case study methodology. 30:26–30.
- Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile. 2017. Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (woebot): a randomized controlled trial. *JMIR mental health*, 4(2).
- Kathleen C. Fraser, Frank Rudzicz, and Graeme Hirst. 2016. [Detecting late-life depression in alzheimer’s disease through analysis of speech and language](#). In *Proceedings of the Third Workshop on CLPsych*. ACL.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in NIPS 29*.
- George Gkotsis, Anika Oellrich, Tim Hubbard, Richard Dobson, Maria Liakata, Sumithra Velupillai, and Rina Dutta. 2016. [The language of mental health problems in social media](#). In *Proceedings of the Third Workshop on CLPsych*. ACL.
- CE Hill, C Greenwald, KG Reed, D Charles, MK O’Farrell, and JA Carter. 1981. Manual for counselor and client verbal response category systems. *Columbus, OH: Marathon*.
- Clara E Hill. 1978. Development of a counselor verbal response category. *Journal of Counseling Psychology*, 25(5):461.
- Clara E Hill, Barbara J Thompson, and Elizabeth Nutt Williams. 1997. A guide to conducting consensual qualitative research. *The counseling psychologist*, 25(4):517–572.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *ACL*. Association for Computational Linguistics.
- Christine Howes, Matthew Purver, and Rose McCabe. 2014. [Linguistic indicators of severity and progress in online text-based therapy for depression](#). In *Proceedings of the Workshop on CLPsych*. ACL.
- Thomas D. Hull. 2015. A preliminary study of talkspaces text-based psychotherapy.
- Zac E Imel, Mark Steyvers, and David C Atkins. 2015. Computational psychotherapy research: Scaling up the evaluation of patient–provider interactions. *Psychotherapy*, 52(1):19.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. Tying word vectors and word classifiers: A loss framework for language modeling. In *ICLR*.
- Zunaira Jamil, Diana Inkpen, Prasadith Buddhitha, and Kenton White. 2017. [Monitoring tweets for depression to detect at-risk users](#). In *Proceedings of the Fourth Workshop on CLPsych*, pages 32–40. ACL.
- Yangfeng Ji, Hwajung Hong, Rosa Arriaga, Agata Rozga, Gregory Abowd, and Jacob Eisenstein. 2014. [Mining themes and interests in the asperger’s and autism community](#). In *Proceedings of the Workshop on CLPsych*. ACL.
- Seung-Hee Jee Hea-Young Oh Kyung-Min Kim. 2010. The response analysis of single session chat counseling - the cases of adolescents with the suicidal crisis -. *The Korea Journal of Youth Counseling*, 18(2):167–186.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on EMNLP*.
- Kate Loveys, Patrick Crutchley, Emily Wyatt, and Glen Coppersmith. 2017. [Small but mighty: Affective micropatterns for quantifying mental health from social media language](#). In *Proceedings of the Fourth Workshop on CLPsych*. ACL.
- Kien Hoa Ly, Ann-Marie Ly, and Gerhard Andersson. 2017. A fully automated conversational agent for promoting mental well-being: A pilot rct using mixed methods. *Internet Interventions*, 10:39–46.
- Akio Mantani, Tadashi Kato, Toshi A Furukawa, Masaru Horikoshi, Hissei Imai, Takahiro Hiroe, Bun Chino, Tadashi Funayama, Naohiro Yonemoto,

- Qi Zhou, et al. 2017. Smartphone cognitive behavioral therapy as an adjunct to pharmacotherapy for refractory depression: randomized controlled trial. *Journal of medical Internet research*, 19(11).
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. Regularizing and optimizing LSTM language models. In *ICLR*.
- Margaret Mitchell, Kristy Hollingshead, and Glen Coppersmith. 2015. [Quantifying the language of schizophrenia in social media](#). In *Proceedings of the 2nd Workshop on CLPsych*. ACL.
- Michelle Morales, Stefan Scherer, and Rivka Levitan. 2017. [A cross-modal review of indicators for depression detection systems](#). In *Proceedings of the Fourth Workshop on CLPsych*. ACL.
- Sylvester Olubolu Orimaye, Jojo Sze-Meng Wong, and Karen Jennifer Golden. 2014. [Learning predictive linguistic features for alzheimers disease and related dementias using verbal utterances](#). In *Proceedings of the Workshop on CLPsych*. ACL.
- Sungjoon Park, Jeongmin Byun, Sion Baek, Yongseok Cho, and Alice Oh. 2018a. Subword-level word vector representations for korean. In *Proceedings of the 56th Annual Meeting of the ACL*.
- Yookoon Park, Jaemin Cho, and Gunhee Kim. 2018b. A hierarchical latent structure for variational conversation modeling.
- Prajit Ramachandran, Peter Liu, and Quoc Le. 2017. Unsupervised pretraining for sequence to sequence learning. In *Proceedings of the 2017 Conference on EMNLP*.
- Philip Resnik, William Armstrong, Leonardo Claudino, Thang Nguyen, Viet-An Nguyen, and Jordan Boyd-Graber. 2015. [Beyond lda: Exploring supervised topic modeling for depression-related language in twitter](#). In *Proceedings of the 2nd Workshop on CLPsych*. ACL.
- Julius Seeman. 1949. A study of the process of nondirective therapy. *Journal of Consulting Psychology*, 13(3):157.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*.
- Judy Hanwen Shen and Frank Rudzicz. 2017. [Detecting anxiety through reddit](#). In *Proceedings of the Fourth Workshop on CLPsych*, pages 58–65. ACL.
- Juan C Torrado, Javier Gomez, and Germán Montoro. 2017. Emotional self-regulation of individuals with autism spectrum disorders: smartwatches for monitoring and interaction. *Sensors*, 17(6):1359.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the NAACL*.
- Andrew Yates, Arman Cohan, and Nazli Goharian. 2017. [Depression and self-harm risk assessment in online forums](#). In *Proceedings of the 2017 Conference on EMNLP*. ACL.

A Appendices

Factual Information	
Counselor	nice to meet you of course first I see. now . . [clients name]s well . .
Client	never visited neuropsychiatry female , worker never received psychological counseling motivation of counseling age # , female
Anecdotal Experience	
Counselor	might be left behind be in a peer relationships I see. actually well . . I understand what you mean
Client	too hard didnt get along seem to be I thought that I was totally wrong
Appealing Problem	
Counselor	right . all but now the relationship is nice to meet you you told me well more comfortable?
Client	sick and sad have many thoughts keep thinking I think I did it no future and frustrating

Psychological Change	
Counselor	right . and if you do that someday suddenly I did I hope so question arises
Client	did not think so but I was so shocked it could have been longer but looking back I think I will
Counseling Process	
Counselor	hard to express but counseling proceeds Hello ? [clients name] , lets talk at # yes the counseling time is
Client	how are you ? what do you think ? thank you for listening hello counselor time is not enough

Table 6: Representative phrases in utterances of both conversational parties, extracted based on the relative importance of the phrase in terms of classification by using attention weights.