

Enhancing Depression Detection via Question-wise Modality Fusion

Aishik Mandal¹, Dana Atzil-Slonim², Thamar Solorio³, Iryna Gurevych¹

¹Ubiquitous Knowledge Processing Lab (UKP Lab)

Department of Computer Science and Hessian Center for AI (hessian.AI)
Technische Universität Darmstadt

²Department of Psychology, Bar-Ilan University

³MBZUAI

www.ukp.tu-darmstadt.de

Abstract

Depression is a highly prevalent and disabling condition that incurs substantial personal and societal costs. Current depression diagnosis involves determining the depression severity of a person through self-reported questionnaires or interviews conducted by clinicians. This often leads to delayed treatment and involves substantial human resources. Thus, several works try to automate the process using multimodal data. However, they usually overlook the following: i) The variable contribution of each modality for each question in the questionnaire and ii) Using ordinal classification for the task. This results in sub-optimal fusion and training methods. In this work, we propose a novel Question-wise Modality Fusion (*QuestMF*) framework trained with a novel Imbalanced Ordinal Log-Loss (*ImbOLL*) function to tackle these issues. The performance of our framework is comparable to the current state-of-the-art models on the E-DAIC dataset and enhances interpretability by predicting scores for each question. This will help clinicians identify an individual's symptoms, allowing them to customise their interventions accordingly. We also make the code¹ for the *QuestMF* framework publicly available.

1 Introduction

Depression is a major cause of disability globally². Its personal and societal impact makes optimising mental health care practices crucial. Existing diagnostic systems of depression are heavily dependent on clinicians' proficiency in attending to patients' verbal and non-verbal cues, but achieving this expertise requires extensive training (Atzil-Slonim et al., 2024). The growing demand for mental health care services, coupled with a shortage of qualified providers, means that many individuals with depression go undiagnosed and untreated

(Kazdin, 2021). Detection of depression severity is crucial, as it can prevent deterioration and enable adequate and effective treatment. Current diagnostic systems have faced criticism for failing to capture the significant heterogeneity and variability of symptoms between individuals (Bickman, 2020). Understanding how different symptoms vary between individuals could lead to more personalised and effective interventions.

A common way to assess depression or track depression throughout a treatment program is based on self-reported questionnaires like PHQ-8 (Kroenke et al., 2009) or BDI-II (Beck et al., 1996). These questionnaires contain questions regarding depression symptoms, and patients need to score each question based on how frequently they encounter these symptoms. The sum of the scores from each question gives the patient's depression severity score. However, such methods burden patients, especially when they are required to complete the questionnaires repeatedly as part of ongoing treatment monitoring (Kazdin, 2008). Thus, to improve the speed and convenience of diagnosing and monitoring depression, efforts are devoted to building depression severity prediction methods through machine learning. Initial works on automatic depression detection focused on using social media data (De Choudhury et al., 2013) for binary depression classification. However, due to the lack of predicting depression severity, such a method is unable to prioritise people with higher levels of depression. So, the task was reformulated as classification among four depression levels (Naseem et al., 2022). However, models trained on social media data are unsuitable for clinical settings (Wongkoblap et al., 2017).

Also, social media data lacks multimodal cues, which are often used by therapists to infer the depression severity of a patient. Depression has identifiable verbal and nonverbal characteristics, such as facial expression (Slonim et al., 2024), prosodic

¹ QuestMF code

²WHO (2023, March 31). Depressive disorder.

information (Cummins et al., 2015; Cohn et al., 2009; Scherer et al., 2014; Paz et al., 2024) and semantic features (Chim et al., 2024). To utilise these cues for depression assessment, the AVEC challenges (Ringeval et al., 2017, 2019) released semi-clinical datasets, DAIC-WOZ (Gratch et al., 2014) and E-DAIC (DeVault et al., 2014), containing recorded interviews and self-reported PHQ-8 questionnaires. These questionnaires help detect symptoms and give more fine-grained depression severity levels.

Various works have tried to utilise the text, audio and video modality from the AVEC datasets and improve the fusion between them (Rodrigues Makiuchi et al., 2019; Sun et al., 2021, 2022; Zhao and Wang, 2022; Ray et al., 2019; Yuan et al., 2024) to predict the depression severity score obtained from PHQ-8 questionnaires. However, these approaches only use one fusion module to fuse text, audio, and video information to predict the depression severity score (the sum of the scores for each question in the questionnaire). This design choice results in a failure to model the variable contributions of each modality depending on the questions in the questionnaire, leading to sub-optimal fusion. For example, a question on being fidgety may require more attention to audio-visual modalities. On the other hand, text transcripts may contribute more significantly to a question regarding a person’s appetite. While Van Steijn et al. (2022) also perform question-wise modeling, they mainly use text features concatenated with a few hand-crafted audio features. So, they do not utilise the audio and video modalities effectively. Another issue is that the current multimodal methods frame the depression severity score prediction as a regression task, resulting in sub-optimal training. Humans score each question in the questionnaire as 0, 1, 2, or 3, depending on the frequency of the symptoms experienced. Thus, the depression severity prediction task should be framed as a question-wise ordinal classification task.

Contributions: We propose a novel Question-wise Modality Fusion (*QuestMF*) framework for depression severity prediction. This framework contains question-wise fusion modules to ensure different contributions from modalities based on the question. In addition, we propose a novel Imbalanced Ordinal Log-Loss (*ImbOLL*) function to train our models for ordinal classification. We find that our method matches the performance of the current state-of-the-art methods on the E-DAIC

dataset and enhances interpretability for clinicians by identifying an individual’s specific symptoms. We also analyse the importance of each modality for each question and find that a fusion of text and video modalities performs best in most questions.

2 Background & Related Work

2.1 Single Modality Methods

Earlier works in depression severity prediction focused on the text modality like the use of linguistic feature extraction based on LIWC (De Choudhury et al., 2013), Bag-of-word models (Nadeem, 2016), word2vec embeddings (Husseini Orabi et al., 2018) or using emotion features (Aragón et al., 2019). With pre-trained language models like BERT (Devlin et al., 2019) improving performances on text-based tasks, depression severity prediction works also utilised them (Rodrigues Makiuchi et al., 2019; Fan et al., 2019; Sun et al., 2022). Van Steijn et al. (2022) introduces a framework to predict scores of each question of a PHQ-8 questionnaire to add interpretability, which is missing in the prior works. These methods also ignore the multi-turn dialogue present in therapy sessions. Thus, Milintsevich et al. (2023) introduces a turn-based method that encodes each dialogue turn using a sentence transformer (Reimers and Gurevych, 2019). We use a similar turn-based model to encode the multi-turn dialogue data in each modality. We, however, use multihead self-attention instead of additive attention to improve the model.

With the advent of LLMs in recent times, Sadeghi et al. (2023) uses GPT-3.5-Turbo³ with encoder models for depression severity prediction. However, its performance falls short of the state-of-the-art models. Moreover, data privacy requirements do not allow data to be sent to proprietary LLMs. These issues motivate us to only explore encoder models.

The AVEC challenges introduced the potential to use audio features for depression detection. This resulted in works utilising low-level audio features (Eyben et al., 2016) extracted by OpenSmile (Eyben et al., 2010). Fan et al. (2019) uses CNNs over the low-level features, while Yin et al. (2019) and Sun et al. (2022) use LSTMs to capture the temporal relation among them. However, LSTMs are sub-optimal at processing long sequences. Thus, Sun et al. (2021) uses transformers to process long

³<https://platform.openai.com/docs/models/gpt-3-5-turbo>

sequences of audio features. There are also other methods that do not use OpenSmile features but rather use spectrograms (Rodrigues Makiuchi et al., 2019) or audio recordings directly (Han et al., 2024; Chen et al., 2023). However, they are computationally expensive, making them difficult to use for multimodal fusion. Here, we use LSTM over low-level features. We break the session into turns and aggregate features at the turn level to make shorter sequences that can be processed using LSTMs.

2.2 Multimodal Fusion Methods

Multimodal methods focus on improving the fusion of the modalities for depression severity prediction. Initial multimodal works used simple concatenation (Rodrigues Makiuchi et al., 2019) or weighted concatenation (Sun et al., 2021) for the fusion of text, audio, and video encodings. Ray et al. (2019) uses attention modules to improve fusion. Some works also use hierarchical fusion at frame level (Yin et al., 2019), word level (Rohanian et al., 2019), or topic level (Guo et al., 2022) to capture the interaction between the modalities at fine-grained levels. For further improvement in fusion, MMFF (Yuan et al., 2024) exploits the high-order interaction between different modalities. However, it is computationally expensive. CubeMLP (Sun et al., 2022) uses MLPs to mix information among modalities to enhance the computational efficiency of fusion. However, it results in lower performance. Zhao and Wang (2022) uses Self-Attention GAN to augment training data to reduce the issue of data shortage. They use a cross-attention based fusion strategy (Tsai et al., 2019). We use the same cross-attention-based fusion. However, these works use a single fusion module, thus ignoring the variable contribution of each modality according to the question. We use question-wise fusion modules to mitigate this issue.

2.3 Ordinal Classification Methods

Ordinal classification has been explored in tasks like sentiment analysis in Twitter (Nakov et al., 2016; Rosenthal et al., 2017). While depression severity score prediction is also an ordinal classification task, prior multimodal methods (Rodrigues Makiuchi et al., 2019; Milintsevich et al., 2023; Ray et al., 2019; Yuan et al., 2024; Zhao and Wang, 2022) treat it as a regression task. As a result, ordinal classification has been rarely explored in depression severity score prediction (Van Steijn et al., 2022). Ordinal classification methods include

ordinal binary classification methods (Frank and Hall, 2001; Allwein et al., 2000), threshold methods (Lin and Li, 2006; Verwaeren et al., 2012; Cao et al., 2020) and loss-sensitive classification methods (Rennie and Srebro, 2005; Diaz and Marathe, 2019; Bertinetto et al., 2020; Cao et al., 2020; Castagnos et al., 2022). However, these methods are not suitable for imbalanced datasets. Since very few patients feel a specific symptom very frequently, the distribution of question-wise scores (labels) in depression severity prediction datasets is imbalanced. Thus, we propose Imbalanced Ordinal Log-Loss (*ImbOLL*), a modified version of the OLL function (Castagnos et al., 2022) to handle the data imbalance

In summary, our work is the first to perform question-wise modality fusion and present a loss function for imbalanced ordinal classification of depression severity prediction task. Moreover, we are the first to analyse the contribution of each modality towards the score of each question, thus improving interpretability.

3 Dataset

Multimodal clinical data collection for depression detection is difficult due to privacy issues, thus resulting in small datasets. In this work, we use the E-DAIC dataset from the AVEC 2019 DDS (Ringeval et al., 2019) challenge. While the E-DAIC dataset is also small, to the best of our knowledge, it is the only dataset with more than 200 data points available for research on depression severity prediction. We do not use the DAIC-WoZ dataset since it is a subset of the E-DAIC dataset. Other available datasets are either smaller (Zou et al., 2023) or are not clinically grounded with self-reported questionnaires (Yoon et al., 2022). The E-DAIC dataset collected recorded interview sessions with a virtual agent and filled out self-reported PHQ-8 questionnaires for each participant. All the interviews were conducted in English. The dataset provides text transcripts of participant dialogues, recorded audio clips, and visual features like ResNet, VGG, and OpenFace for each interview session. The recorded videos have not been released due to privacy concerns. The dataset contains 275 sessions. The training set includes 163 sessions, and the validation and test sets each contain 56 sessions. However, one session from the validation and one from the test set have incomplete video feature files. Thus, we do not use them in the evaluation. The dataset

also provides the PHQ-8 scores of all participants. The PHQ-8 score ranges from 0 to 24. While the training and validation sets contain scores for each of the eight PHQ-8 questions (0 to 3), the test split only contains the total PHQ-8 scores. More details on the PHQ-8 questionnaire are provided in Appendix A.

4 QuestMF Framework

In this section, we present our novel Question-wise Modality Fusion (*QuestMF*) framework. In this framework, we use n different single modality encoders for each modality and n different modality fusion models corresponding to n questions (thus question-wise modality fusion) in a questionnaire. Each of the n fused models outputs the score for its corresponding question (0, 1, 2, or 3), ensuring different contributions from each modality depending on the question, which was lacking in previous works. These question-wise scores are added to get the total questionnaire score (0 to $3n$). Figure 1 shows the proposed framework. Moreover, the *QuestMF* framework improves interpretability by predicting the score of each question. This allows clinicians to understand the symptoms affecting a patient and create interventions accordingly. This framework will also enable clinicians to track the progression of each symptom through the question-wise scores throughout the multiple therapy sessions during the treatment.

Moreover, current multimodal methods (Rodrigues Makiuchi et al., 2019; Milintsevich et al., 2023; Ray et al., 2019; Yuan et al., 2024; Zhao and Wang, 2022) train the multimodal methods to predict the total depression severity score as a regression task treating it as a continuous variable. However, the question-wise scores belong to 4 classes: 0, 1, 2, 3 depending on the frequency of the symptoms experienced, thus making it an ordinal classification task. Additionally, treating question-wise scores as continuous variables also reduces interpretability as fractional scores like 1.5 can mean experiencing a symptom at the frequency of either score 1 or score 2. Thus, framing the question-wise score prediction as an ordinal classification task ensures improved interpretability as we get the predicted probabilities of the 4 classes: 0, 1, 2, 3 and choose the class with the highest probability. This is also more similar to how humans fill out the questionnaires.

Next, we discuss the single modality encoders

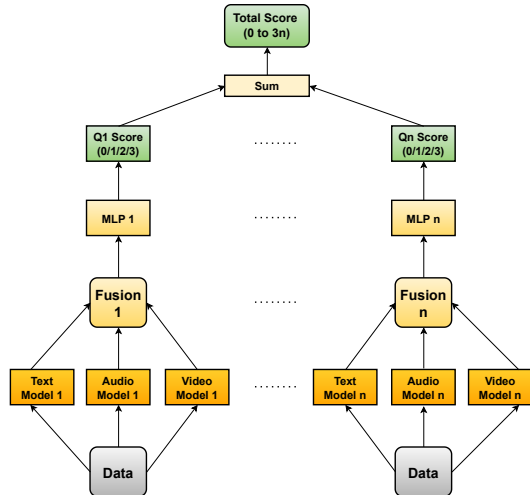


Figure 1: Proposed *QuestMF* framework to predict depression severity score. Here, Q_x denotes Question number x in the questionnaire. MLP denotes Multilayer Perceptron, which is used as the classification head. Each question is scored among classes $\{0, 1, 2, 3\}$. These scores are then added to get the total score $\in \{0, 1, 2, \dots, 3n\}$.

used in the framework in Section 4.1 and the fusion methods used to combine the single modality encodings in Section 4.2. Finally, we introduce the novel *ImbOLL* function used to train the models for ordinal classification in Section 4.3.

4.1 Single Modality Encoder Models

All the single modality encoder models follow a turn-based method similar to Milintsevich et al. (2023) to better encode the interviews containing multi-turn dialogues. We start by encoding the dialogue turns and then use these turn encodings to generate an encoding for the whole session. The overall structure is uniform across all modalities, as shown in Figure 2. Now, we describe each single modality encoder in detail.

4.1.1 Text Encoder Model

For text, we use the textual transcripts from interview sessions. We break the transcripts into dialogue turns and only consider the dialogue turns from participants. We encode the turns using a pre-trained sentence transformer (Reimers and Gurevych, 2019). Now, we get turn encoding, $X_i \in \mathbb{R}^{tokens_i \times D_{model}}$, where $i \in (1, 2, \dots, m)$ for m turns, $tokens_i$ is the number of tokens in turn i and D_{model} is the model output dimension. Next, we use mean pooling over the tokens and normalise them to get $X_{mean,i} \in \mathbb{R}^{D_{model}}$ (following Reimers and Gurevych (2019)). Then, we

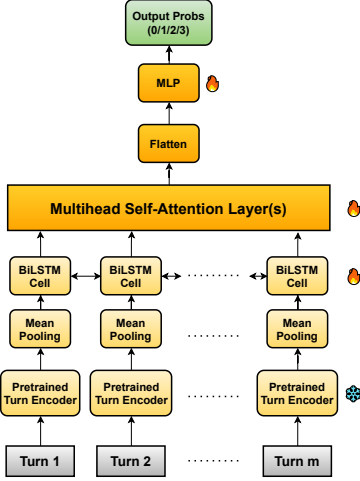


Figure 2: Architecture of single modality encoder models. We use a turn-based architecture to encode multi-turn dialogue data.

pass these turn encodings through a Bidirectional LSTM layer. The Bidirectional LSTM layer ensures that the turns can interact among themselves. This gives us $X_{lstm,i} \in \mathbb{R}^{2 \cdot D_{lstm}}$. Next, we use a multihead attention layer to determine the importance of each turn and get an updated encoding $X_{att,i} \in \mathbb{R}^{2 \cdot D_{lstm}}$. We do not add positional embeddings to $X_{lstm,i}$ as the attention layer is only used to get the importance of each turn based on only the turn contents. Next, we concatenate and flatten the turn encodings obtained after the multihead attention layer to get a session-level encoding representation of $X_{session} \in \mathbb{R}^{(2m \cdot D_{lstm})}$. Finally, we pass this session-level encoding through a multilayer perceptron (MLP) to get the score probabilities.

4.1.2 Audio Encoder Model

For audio, we use the low-level features (a set of basic acoustic parameters used for voice research and affective computing suggested by Eyben et al. (2016)) extracted by OpenSmile. Details on the low-level features extracted using OpenSmile are provided in Appendix C. These features are extracted at every 0.01 seconds. Like the text model, we process the information at the turn level. For each dialogue turn i , the dataset contains a starting time $t_{start,i}$ and an ending time $t_{end,i}$. We get the features extracted from time $t_{start,i}$ to $t_{end,i}$ and apply mean pooling to get the aggregated features in a turn, $X_{mean,i} \in \mathbb{R}^{D_{Features}}$. Here, $D_{Features}$ is the number of the features extracted by OpenSmile. After this, we pass them through a Bi-LSTM layer and two attention layers to get updated turn encod-

ings. These are concatenated, flattened, and passed through an MLP to get the score probabilities.

4.1.3 Video Encoder Model

For the video encoder model, we use ResNet (He et al., 2016) features. Similar to text and audio, we aggregate information at turn level. For this, we get the ResNet features for the frames in a dialogue turn, i.e., frames in $t_{start,i} * sr : t_{end,i} * sr$, where $t_{start,i}$ is the starting time and $t_{end,i}$ is the ending time of the dialogue turn and sr is the frames rate at which the video is recorded. We pass these ResNet features in a dialogue turn through a mean pooling layer and normalise them to get turn-level encoding $X_{mean,i} \in \mathbb{R}^{D_{ResNet}}$ where D_{ResNet} is the output dimension of the ResNet model. After this, we follow the same architecture as the audio encoder to get the score probabilities.

4.2 Modality Fused Models

For modality fusion, we use cross-attention based late fusion method introduced by Tsai et al. (2019). We do not explore early fusion since late fusion is shown to be better in fusion of text, audio and video modality (Snoek et al., 2005) and also more commonly used over early fusion. The cross-attention layers are sometimes accompanied by $X \rightarrow Y$. This denotes that the encoding of the Y modality is used as the query, and the encoding of the X modality is used as the key and value in the cross-attention layer. Next, we describe the modality fused models in detail.

4.2.1 Two-Modality Fused Models

Figure 3 shows our two-modality fused models. We use the output from the multihead attention layers of the trained single modality encoders as the modality encoding. We use multihead cross-attention layers over these encodings to exchange information among the modalities. Considering modality encodings $M1$ and $M2$, we use two cross-attention layers $M1 \rightarrow M2$ and $M2 \rightarrow M1$ for interaction among the modalities. This is followed by a multihead self-attention layer for each cross-attention layer. Finally, we concatenate the encodings from the self-attention layers to get a fused encoding. We flatten this fused encoding and pass it through an MLP to get the score probabilities.

4.2.2 Three-Modality Fused Models

Figure 4 shows our three-modality fused model. We use the output from the multihead attention layers of the trained single modality encoders as

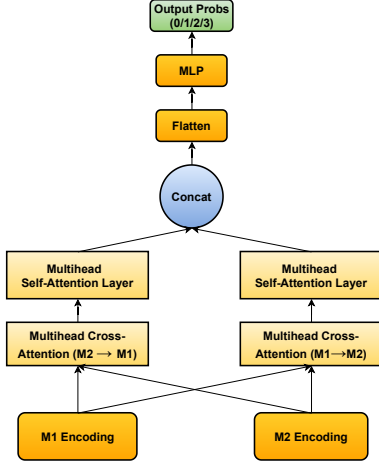


Figure 3: Architecture of two-modality fused models. We use cross-attention layers for interaction among modalities $M1$ and $M2$. In cross-attention, $X \rightarrow Y$ denotes that the Y modality encoding is used as the query and the X modality encoding as the key and value.

the modality encoding. Then, we use multihead cross-attention layers to pass information among the modalities. In this case, we have six combinations of query and (key, value) pairs. Now, we accumulate the encoding for each modality with information from the other two modalities. We perform this by concatenating the outputs from two cross-attention layers using the same modality as the query. For example, the audio encoding with information from text and video modalities will concatenate encodings obtained from cross-attention layers $T \rightarrow A$ and $V \rightarrow A$. We pass these encodings for each modality with the information from other modalities through a multihead self-attention layer. Next, we concatenate them to get a combined encoding of all three modalities. Finally, we flatten the combined encoding and pass it through an MLP to get the score probabilities.

4.3 ImbOLL Function

Now, we introduce the novel *ImbOLL* function we use to train our models. The *ImbOLL* function is a modified version of the OLL (Castagnos et al., 2022) function. The OLL function is used to train models for ordinal classification. The OLL function for N classes is defined as:

$$\mathcal{L}_{OLL-\alpha}(P, y) = - \sum_{i=1}^N \log(1 - p_i) d(y, i)^\alpha \quad (1)$$

where y is the actual class, p_i is the predicted probability of class i , α is a hyperparameter and

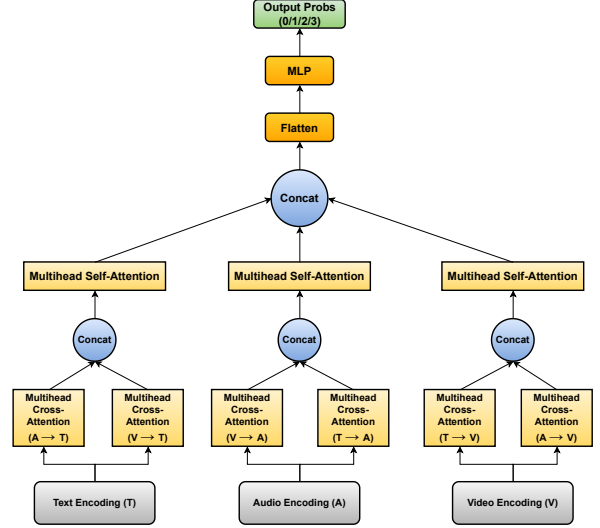


Figure 4: Architecture of the three-modality fused model. In cross-attention, $X \rightarrow Y$ denotes that the Y modality encoding is used as the query and the X modality encoding as the key and value.

$d(y, i)$ is the distance between the classes y and i which is defined as:

$$d(y, i) = |y - i| \quad (2)$$

The OLL function is based on the principle of penalising a model for bad decisions instead of rewarding good decisions. However, the OLL function is not suitable for imbalanced datasets. The questions of the PHQ-8 questionnaire consist of 4 possible classes according to the frequency of symptoms: 0, 1, 2, and 3. However, very few participants give a score of 2 or 3 to a question as very few patients feel a particular symptom so often. This results in an imbalanced score distribution. To consider this, we introduce weights $w(y)$, which gives a harsher punishment to a model when it makes a wrong decision for a rarer ground truth score. This is inspired by the use of weighted cross-entropy loss in imbalanced classification tasks (Aurelio et al., 2019). The weights are defined as:

$$w(y) = \frac{n_T}{n_y} \quad (3)$$

Where n_T is the total number of data points in the training set, and n_y is the number of data points in the training set belonging to class y . Our novel loss function, *ImbOLL*, is defined as follows:

$$\mathcal{L}_{ImbOLL-\alpha, \beta} = - \sum_{i=1}^N \log(1 - p_i) d(y, i)^\alpha w(y)^\beta \quad (4)$$

Model	Modalities	CCC(↑)	RMSE(↓)	MAE(↓)
Ray et al. (2019)	Text, Audio, Video	0.67	4.73	4.02
Sun et al. (2022)	Text, Audio, Video	0.583	-	4.37
Zhao and Wang (2022)	Text, Audio, Video	-	4.14	3.56
Yuan et al. (2024)	Text, Audio, Video	0.676	4.91	3.98
Van Steijn et al. (2022)	Text, Audio	0.62	6.06	-
<i>Total</i>	Text, Audio, Video	0.618	4.99	4.03
<i>QuestMF</i> (MSE)	Text, Audio, Video	0.620	5.31	4.16
<i>QuestMF</i> (OLL)	Text, Audio, Video	0.656	5.17	3.89
<i>QuestMF</i> (<i>ImbOLL</i>)	Text, Audio, Video	0.685	5.32	4.11

Table 1: Results of *QuestMF* trained with *ImbOLL* function compared with ablation frameworks and prior works.

where α and β are hyperparameters.

5 Experiments

In this section, we describe the experiments with *QuestMF* and its ablation frameworks. We use the following ablation frameworks:

Total: We train the models with the MSE loss function to predict the total questionnaire score. This framework consists of a single modality encoder for each modality and a single fused model.

***QuestMF* (MSE):** We train the *QuestMF* framework with the MSE loss function.

***QuestMF* (OLL):** We train the *QuestMF* framework with OLL function.

***QuestMF* (*ImbOLL*):** We train the *QuestMF* framework with *ImbOLL* function. *This is our proposed framework.*

To evaluate the performance of the methods in the depression severity prediction task, we use the standard metrics used in prior works: Concordance Correlation Coefficient (CCC), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE). CCC is defined as:

$$\rho_c = \frac{2\rho\sigma_x\sigma_y}{\sigma_x^2 + \sigma_y^2 + (\mu_x - \mu_y)^2} \quad (5)$$

Where ρ is the Pearson correlation between variables x and y . σ_x and σ_y are the standard deviations of variables x and y . μ_x and μ_y are the means of variables x and y . We use CCC as the primary metric because it is unbiased by changes in scale and location (Lin, 1989). Psychologists use CCC to assess the agreement between test scores from different raters. It was also used as the evaluation metric for the AVEC 2019 challenge. We also report RMSE and MAE. A higher CCC score is desirable to show that the predicted and actual outputs

correlate well. For RMSE and MAE, a lower value is desired, as it shows a smaller difference between the predicted and the actual output. More details on the metrics are provided in Appendix B. For the *ImbOLL* function, we empirically find the parameters $\alpha = 1$ and $\beta = 0.5$ to be the best. For the OLL function, we find $\alpha = 1$ to give the best results. The detailed results of the experiments with hyperparameters of *ImbOLL* and OLL are presented in Appendix D. To show the robustness of our model, we run our experiments on three different seeds: 42, 100, and 1234. We are the first in this domain to run experiments on multiple seeds. The training strategy, checkpoint selection, and hyperparameter details of the models are provided in Appendix E.

6 Results & Analysis

The results comparing our proposed *QuestMF* (*ImbOLL*) framework with its ablations and current state-of-the-art methods are presented in Table 1. Since the prior works only show their best results on a single run, we also pick our best results on CCC for a fair comparison. As we can see, our proposed *QuestMF* (*ImbOLL*) framework matches the performance of state-of-the-art models in CCC, the primary metric used to evaluate depression severity prediction tasks. In addition to comparable performance, *QuestMF* framework provides questionnaire scores, improving interpretability over current methods, thus allowing clinicians to design personalised interventions. We also show the robustness of our frameworks over multiple runs, which is not done by the previous works. The mean and standard deviation of the performance over 3 runs are presented in Table 2. We observe that *Total* and *QuestMF* (MSE) perform similarly. However, training with ordinal classification objective improves

Model	CCC(\uparrow)	RMSE(\downarrow)	MAE(\downarrow)
<i>Total</i>	0.602 ± 0.015	5.10 ± 0.08	3.99 ± 0.05
<i>QuestMF</i> (MSE)	0.602 ± 0.024	5.36 ± 0.24	4.21 ± 0.12
<i>QuestMF</i> (OLL)	0.640 ± 0.018	5.14 ± 0.05	3.88 ± 0.01
<i>QuestMF</i> (<i>ImbOLL</i>)	0.662 ± 0.022	5.25 ± 0.08	3.95 ± 0.13

Table 2: Results of *QuestMF* (*ImbOLL*) framework over 3 different seed runs compared with ablation frameworks.

the performance, as we can see from the results of *QuestMF* (OLL). This shows the effectiveness of combining *QuestMF* with ordinal classification training. Training with our novel *ImbOLL* function further improves the results on CCC. It is also robust, with a standard deviation of 0.022.

Next, we present an ablation study to observe the performance of single modality models, two-modality fused models, and the three-modality fused model with the different frameworks. We present the results in Table 3. For all single modality models and two-modality fused models, we observe that *QuestMF* (OLL) and *QuestMF* (*ImbOLL*) frameworks show better performance than the other frameworks. All the frameworks except *QuestMF* (*ImbOLL*) and *QuestMF* (OLL) show the best performance with only text and video fusion, while the performance drops when all three modalities are fused. This shows that training with an ordinal classification task gives a better optimisation objective. Moreover, we also see that *QuestMF* (*ImbOLL*) and *QuestMF* (OLL) show the best performance gains when adding more modalities. Comparing the results of Text + Audio + Video with the Text models, we observe that *QuestMF* (*ImbOLL*) achieves an improvement of 0.047 (Text + Audio + Video (0.662) – Text (0.615)) on CCC and *QuestMF* (OLL) achieves an improvement of 0.048 (Text + Audio + Video (0.640) – Text (0.592)). Among the regression methods, *Total* achieves the best improvement when comparing the results of Text + Audio + Video with the Text models. However, it only achieves an improvement of 0.011 (Text + Audio + Video (0.602) – Text(0.591)). This further shows that the *QuestMF* (*ImbOLL*) framework improves fusion. Appendix F provides more detailed results for this ablation.

Finally, we analyse the importance of each modality toward predicting the score for each question in the framework. Since we lack fine-grained labels in the test split, we use the validation set CCC for each question in this analysis. A higher

Modalities	Model	CCC(\uparrow)
T	<i>Total</i>	0.591
	<i>QuestMF</i> (MSE)	0.593
	<i>QuestMF</i> (OLL)	0.592
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.615
A	<i>Total</i>	0.212
	<i>QuestMF</i> (MSE)	0.239
	<i>QuestMF</i> (OLL)	0.264
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.273
V	<i>Total</i>	-0.067
	<i>QuestMF</i> (MSE)	-0.075
	<i>QuestMF</i> (OLL)	-0.041
	<i>QuestMF</i> (<i>ImbOLL</i>)	-0.052
T+A	<i>Total</i>	0.607
	<i>QuestMF</i> (MSE)	0.618
	<i>QuestMF</i> (OLL)	0.628
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.643
T+V	<i>Total</i>	0.610
	<i>QuestMF</i> (MSE)	0.627
	<i>QuestMF</i> (OLL)	0.628
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.659
A+V	<i>Total</i>	0.058
	<i>QuestMF</i> (MSE)	0.070
	<i>QuestMF</i> (OLL)	0.139
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.159
T+A+V	<i>Total</i>	0.602
	<i>QuestMF</i> (MSE)	0.602
	<i>QuestMF</i> (OLL)	0.640
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.662

Table 3: Ablation results for using different combinations of modalities with different frameworks. Here, T refers to Text, A refers to Audio, and V refers to Video. An addition between the modalities denotes using a fusion of them. The CCC scores presented are the mean over 3 different seed runs.

CCC shows greater importance. The CCC for each question with different modality combinations is shown in Figure 5. From them, we observe:

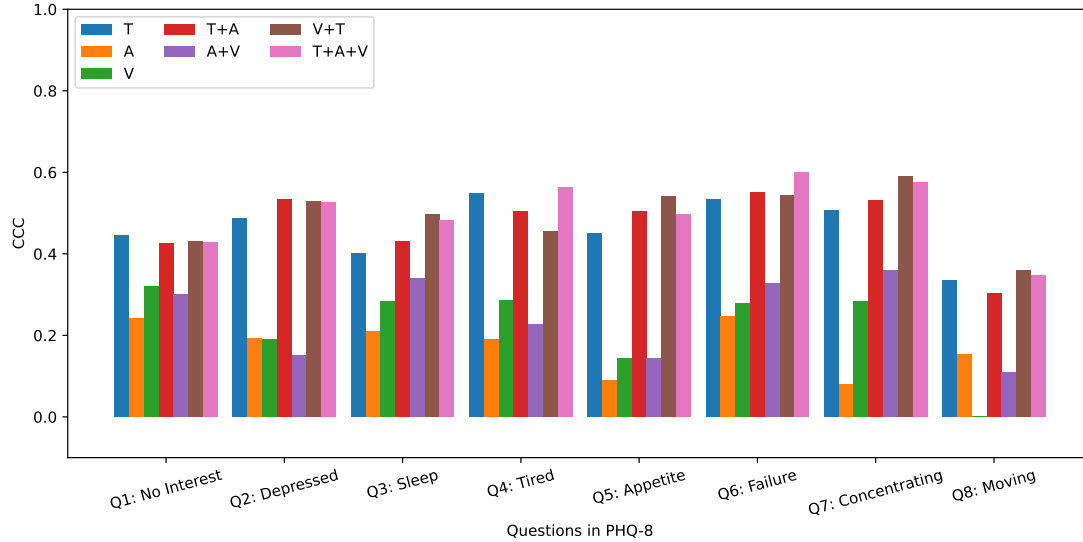


Figure 5: Validation CCC for each question with different modality models. Here, T refers to Text, A refers to Audio, and V refers to Video. An addition between the modalities denotes using a fusion of them. The video model for question 8 gives the same scores to all data points. Thus, its CCC is not valid and is not shown in the graph.

Q1: Feeling no interest. Text modality gives the best results. This may be because the content of the conversation in the interview is the best indicator to determine loss of interest in hobbies.

Q2: Feeling depressed. Text + Audio gives the best results with very close results from Text + Video and Text + Audio + Video. Audio cues like flat speech and visual cues like sadness or blunted facial expressions might help predict this score. However, a fusion of all three struggles to train due to a small number of training data.

Q3: Irregular sleep. Text + Video gives the best results. Sleep disturbances are often directly reported by individuals, which explains the contribution of the text modality. The video modality contributes since sleep issues can often be observed from posture and general demeanour.

Q4: Feeling tired. Text + Audio + Video gives the best results as tiredness is often visible in a person’s face, body language, and way of speaking.

Q5: Irregular appetite. Text + Video gives the best results. A person’s appetite can be determined by directly asking them, so the text has the highest importance here. The video also contributes as physical appearance may influence the prediction.

Q6: Feeling like a failure. Text + Audio + Video gives the best results. Acoustic cues like a disappointed voice and visual cues like a saddened face help in the prediction.

Q7: Trouble concentrating. Text + video gives the best results. Visual cues like gaze can be an important factor. Looking away and not making

eye contact may indicate concentration problems.

Q8: Irregularities in moving and speaking. In this case, Text + Video gives the best results, with Text + Audio + Video closely following. Movement can be captured through the video, and irregularities in a speech can be captured from transcripts and audio recordings.

7 Conclusions & Future Work

In this work, we show that our Question-wise Modality Fusion (*QuestMF*) framework trained with Imbalanced Ordinal Log-Loss (*ImbOLL*) function improves the interpretability in depression severity score prediction by predicting scores of specific questions. This can help clinicians identify particular symptoms or symptom combinations, enabling them to tailor their interventions to the individual’s specific needs. The *QuestMF* (*ImbOLL*) framework also shows performance comparable to current state-of-the-art models on the E-DAIC dataset. We also show its robustness over different seeds. Our framework can assist clinicians in diagnosing and monitoring depression and reduce the burden placed on patients in filling out self-reported questionnaires. Additionally, we perform an extensive analysis to understand the importance of each modality for each question in the questionnaire. By releasing the code, we hope to enable future research of this framework on other questionnaires for mental health assessment and on real world longitudinal therapy data.

Limitations

While the question-wise modality fusion framework trained with *ImbOLL* function offers a solution to considering the variable contribution from modalities based on questions and framing the problem as an ordinal classification task, the data used for training and evaluation are not ideal. While the E-DAIC dataset was released to improve multimodal research in depression severity prediction, the training split only contains 163 sessions. As a result, the trained models are prone to overfitting and high bias and are unlikely to perform well in out-of-distribution data. The validation and test splits also contain only 56 sessions each. As a result, they are far from representing the general population. Moreover, bigger and more diverse datasets are unavailable due to privacy issues. Thus, the *QuestMF (ImbOLL)* framework is only tested on the E-DAIC dataset, which further constrains testing the generalisability of the model. Thus, the results and analysis obtained in this work need to be verified with bigger and more diverse datasets in the future. Also, the E-DAIC dataset only contains first time interviews of participants with a virtual agent similar to enrolling interviews for therapy and does not contain real therapy session interviews. Real world depression tracking also requires longitudinal data, i.e., multiple therapy sessions with the same participant and tracking changes in depression severity throughout their treatment. Since the E-DAIC dataset does not contain such data, we cannot test the effectiveness of our model in such real world situations.

Another limitation is the language and culture coverage. In this work, we only cover the English language, and the dataset is collected in the US. However, people use different languages to express themselves, and people from different cultures express themselves differently, thus affecting therapy. However, *QuestMF (ImbOLL)* could not be developed and tested for such generalisation due to the lack of suitable datasets.

Our focus in this work is to present a more intuitive methodology that considers the variable contribution from modalities according to the question in a questionnaire, frames the task in its true nature of ordinal classification task, gives question-wise scores that can help clinicians design more personalised interventions and analyse the results to understand the contribution of each modality towards the score of each question.

Ethical Considerations

While this work is focused on presenting a methodology and analysis for automatic depression detection, the methods need to be trained on larger datasets to ensure the method's generalisation capabilities. The method should also be assessed for generalisability through clinical trials. Deploying these methods without proper training and assessment through clinical trials could lead to introducing harmful biases in real world situation. Therefore, the framework trained with the E-DAIC dataset may not be used in clinical practice. It requires a broader evaluation and clinical validation before being used in real-world clinical settings.

Acknowledgements

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. This work has also been supported by the DYNAMIC center, which is funded by the LOEWE program of the Hessian Ministry of Science and Arts (Grant Number: LOEWE/1/16/519/03/09.001(0009)/98).

References

- Erin L Allwein, Robert E Schapire, and Yoram Singer. 2000. [Reducing multiclass to binary: A unifying approach for margin classifiers](#). *Journal of machine learning research*, 1(Dec):113–141.
- Mario Ezra Aragón, Adrian Pastor López-Monroy, Luis Carlos González-Gurrola, and Manuel Montes-y Gómez. 2019. [Detecting depression in social media using fine-grained emotions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1481–1486, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dana Atzil-Slonim, Juan Martin Gomez Penedo, and Wolfgang Lutz. 2024. [Leveraging novel technologies and artificial intelligence to advance practice-oriented research](#). *Administration and Policy in Mental Health and Mental Health Services Research*, 51(3):306–317.
- Yuri Sousa Aurelio, Gustavo Matheus de Almeida, Cristiano Leite Castro, and Antônio de Pádua Braga. 2019. [Learning from imbalanced data sets with weighted cross-entropy function](#). *Neural Process. Lett.*, 50(2):1937–1949.

- Aaron T Beck, Robert A Steer, and Gregory K Brown. 1996. [Beck depression inventory](#). San Antonio, TX: The Psychological Corporation.
- Luca Bertinetto, Romain Müller, Konstantinos Tertikas, Sina Samangooei, and Nicholas A. Lord. 2020. [Making better mistakes: Leveraging class hierarchies with deep networks](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 12503–12512. Computer Vision Foundation / IEEE.
- Leonard Bickman. 2020. [Improving mental health services: A 50-year journey from randomized experiments to artificial intelligence and precision mental health](#). *Administration and Policy in Mental Health and Mental Health Services Research*, 47(5):795–843.
- Wenzhi Cao, Vahid Mirjalili, and Sebastian Raschka. 2020. [Rank consistent ordinal regression for neural networks with application to age estimation](#). *Pattern Recognition Letters*, 140:325–331.
- François Castagnos, Martin Mihelich, and Charles Dognin. 2022. [A simple log-based loss function for ordinal text classification](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4604–4609, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Weidong Chen, Xiaofen Xing, Xiangmin Xu, Jianxin Pang, and Lan Du. 2023. [Speechformer++: A hierarchical efficient framework for paralinguistic speech processing](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:775–788.
- Jenny Chim, Adam Tsakalidis, Dimitris Gkoumas, Dana Atzil-Slonim, Yaakov Ophir, Ayah Zirikly, Philip Resnik, and Maria Liakata. 2024. [Overview of the CLPsych 2024 shared task: Leveraging large language models to identify evidence of suicidality risk in online posts](#). In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, pages 177–190, St. Julians, Malta. Association for Computational Linguistics.
- Jeffrey F. Cohn, Tomas Simon Kruez, Iain Matthews, Ying Yang, Minh Hoai Nguyen, Margara Tejera Padilla, Feng Zhou, and Fernando De la Torre. 2009. [Detecting depression from facial actions and vocal prosody](#). In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*, pages 1–7.
- Nicholas Cummins, Stefan Scherer, Jarek Krajewski, Sebastian Schnieder, Julien Epps, and Thomas F Quatieri. 2015. [A review of depression and suicide risk assessment using speech analysis](#). *Speech Communication*, 71(C):10–49.
- Munmun De Choudhury, Scott Counts, and Eric Horvitz. 2013. [Social media as a measurement tool of depression in populations](#). In *Proceedings of the 5th Annual ACM Web Science Conference, WebSci '13*, page 47–56, New York, NY, USA. Association for Computing Machinery.
- David DeVault, Ron Artstein, Grace Benn, Teresa Dey, Edward Fast, Alesia Gainer, Kallirroi Georgila, Jonathan Gratch, Arno Hartholt, Margaux Lhommet, Gale M. Lucas, Stacy Marsella, Fabrizio Morbini, Angela Nazarian, Stefan Scherer, Giota Stratou, Apar Suri, David R. Traum, Rachel Wood, Yuyu Xu, Albert A. Rizzo, and Louis-Philippe Morency. 2014. [Simsensei kiosk: a virtual human interviewer for healthcare decision support](#). In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14, Paris, France, May 5-9, 2014*, pages 1061–1068. IFAAMAS/ACM.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Raul Diaz and Amit Marathe. 2019. [Soft labels for ordinal regression](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4738–4747. Computer Vision Foundation / IEEE.
- Florian Eyben, Klaus R. Scherer, Björn W. Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y. Devillers, Julien Epps, Petri Laukka, Shrikanth S. Narayanan, and Khiet P. Truong. 2016. [The geneva minimalistic acoustic parameter set \(gemaps\) for voice research and affective computing](#). *IEEE Trans. Affect. Comput.*, 7(2):190–202.
- Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. [Opensmile: the munich versatile and fast open-source audio feature extractor](#). In *Proceedings of the 18th ACM International Conference on Multimedia, MM '10*, page 1459–1462, New York, NY, USA. Association for Computing Machinery.
- Weiquan Fan, Zhiwei He, Xiaofen Xing, Bolun Cai, and Weirui Lu. 2019. [Multi-modality depression detection via multi-scale temporal dilated cnns](#). In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop, AVEC '19*, page 73–80, New York, NY, USA. Association for Computing Machinery.
- Eibe Frank and Mark Hall. 2001. [A simple approach to ordinal classification](#). In *Machine Learning: ECML 2001*, pages 145–156, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Jonathan Gratch, Ron Artstein, Gale M Lucas, Giota Stratou, Stefan Scherer, Angela Nazarian, Rachel Wood, Jill Boberg, David DeVault, Stacy Marsella, et al. 2014. [The distress analysis interview corpus of human and computer interviews](#). In *LREC*, pages 3123–3128. Reykjavik.

- Yanrong Guo, Chenyang Zhu, Shijie Hao, and Richang Hong. 2022. [A topic-attentive transformer-based model for multimodal depression detection](#). *CoRR*, abs/2206.13256.
- Zhuojin Han, Yuanyuan Shang, Zhuhong Shao, Jingyi Liu, Guodong Guo, Tie Liu, Hui Ding, and Qiang Hu. 2024. [Spatial-temporal feature network for speech-based depression recognition](#). *IEEE Trans. Cogn. Dev. Syst.*, 16(1):308–318.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Ahmed Husseini Orabi, Prasadith Buddhitha, Mahmoud Husseini Orabi, and Diana Inkpen. 2018. [Deep learning for depression detection of Twitter users](#). In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 88–97, New Orleans, LA. Association for Computational Linguistics.
- Alan E Kazdin. 2008. [Evidence-based treatment and practice: new opportunities to bridge clinical research and practice, enhance the knowledge base, and improve patient care](#). *American psychologist*, 63(3):146.
- Alan E Kazdin. 2021. [Extending the scalability and reach of psychosocial interventions](#). *Bergin and Garfield's handbook of psychotherapy and behavior change: 50th anniversary edition*, pages 763–789.
- Kurt Kroenke, Tara W Strine, Robert L Spitzer, Janet BW Williams, Joyce T Berry, and Ali H Mokdad. 2009. [The phq-8 as a measure of current depression in the general population](#). *Journal of affective disorders*, 114(1-3):163–173.
- Hsuan-Tien Lin and Ling Li. 2006. [Large-margin thresholded ensembles for ordinal regression: Theory and practice](#). In *Algorithmic Learning Theory*, pages 319–333, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Lawrence I-Kuei Lin. 1989. [A concordance correlation coefficient to evaluate reproducibility](#). *Biometrics*, 45(1):255–268.
- Kirill Milintsevich, Kairit Sirts, and Gaël Dias. 2023. [Towards automatic text-based estimation of depression through symptom prediction](#). *Brain Informatics*, 10(1):4.
- Moin Nadeem. 2016. [Identifying depression on twitter](#). *CoRR*, abs/1607.07384.
- Preslav Nakov, Alan Ritter, Sara Rosenthal, Fabrizio Sebastiani, and Veselin Stoyanov. 2016. [Semeval-2016 task 4: Sentiment analysis in twitter](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*, pages 1–18. The Association for Computer Linguistics.
- Usman Naseem, Adam G. Dunn, Jinman Kim, and Matloob Khushi. 2022. [Early identification of depression severity levels on reddit using ordinal classification](#). In *Proceedings of the ACM Web Conference 2022, WWW '22*, page 2563–2572, New York, NY, USA. Association for Computing Machinery.
- Adar Paz, Eshkol Rafaeli, Eran Bar-Kalifa, Eva Gilboa-Schechtman, Sharon Gannot, Shrikanth S Narayanan, and Dana Atzil-Slonim. 2024. [Multimodal analysis of temporal affective variability within treatment for depression](#). *Journal of Consulting and Clinical Psychology*, 92(9):569–581.
- Anupama Ray, Siddharth Kumar, Rutvik Reddy, Pre-rana Mukherjee, and Ritu Garg. 2019. [Multi-level attention network using text, audio and video for depression prediction](#). In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop, AVEC '19*, page 81–88, New York, NY, USA. Association for Computing Machinery.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Jason DM Rennie and Nathan Srebro. 2005. [Loss functions for preference levels: Regression with discrete ordered labels](#). In *Proceedings of the IJCAI multidisciplinary workshop on advances in preference handling*, volume 1. AAAI Press, Menlo Park, CA.
- Fabien Ringeval, Björn Schuller, Michel Valstar, Nicholas Cummins, Roddy Cowie, Leili Tavabi, Maximilian Schmitt, Sina Alisamir, Shahin Amiriparian, Eva-Maria Messner, Siyang Song, Shuo Liu, Ziping Zhao, Adria Mallol-Ragolta, Zhao Ren, Mohammad Soleymani, and Maja Pantic. 2019. [Avec 2019 workshop and challenge: State-of-mind, detecting depression with ai, and cross-cultural affect recognition](#). In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop, AVEC '19*, page 3–12, New York, NY, USA. Association for Computing Machinery.
- Fabien Ringeval, Björn Schuller, Michel Valstar, Jonathan Gratch, Roddy Cowie, Stefan Scherer, Sharon Mozgai, Nicholas Cummins, Maximilian Schmitt, and Maja Pantic. 2017. [Avec 2017: Real-life depression, and affect recognition workshop and challenge](#). In *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge, AVEC '17*, page 3–9, New York, NY, USA. Association for Computing Machinery.
- Mariana Rodrigues Makiuchi, Tifani Warnita, Kuniaki Uto, and Koichi Shinoda. 2019. [Multimodal](#)

- fusion of bert-cnn and gated cnn representations for depression detection. In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop, AVEC '19*, page 55–63, New York, NY, USA. Association for Computing Machinery.
- Morteza Rohanian, Julian Hough, and Matthew Purver. 2019. Detecting depression with word-level multimodal fusion. In *20th Annual Conference of the International Speech Communication Association, Interspeech 2019, Graz, Austria, September 15-19, 2019*, pages 1443–1447. ISCA.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 502–518, Vancouver, Canada. Association for Computational Linguistics.
- Misha Sadeghi, Bernhard Egger, Reza Agahi, Robert Richer, Klara Capito, Lydia Helene Rupp, Lena Schindler-Gmelch, Matthias Berking, and Bjoern M. Eskofier. 2023. Exploring the capabilities of a language model-only approach for depression detection in text data. In *IEEE EMBS International Conference on Biomedical and Health Informatics, BHI 2023, Pittsburgh, PA, USA, October 15-18, 2023*, pages 1–5. IEEE.
- Stefan Scherer, Giota Stratou, Gale M. Lucas, Marwa Mahmoud, Jill Boberg, Jonathan Gratch, Albert A. Rizzo, and Louis-Philippe Morency. 2014. Automatic audiovisual behavior descriptors for psychological disorder analysis. *Image Vis. Comput.*, 32(10):648–658.
- Dana Atzil Slonim, Ido Yehezkel, Adar Paz, Eran Bar-Kalifa, Maya Wolff, Avinoam Dar, and Eva Gilboa-Schechtman. 2024. Facing change: using automated facial expression analysis to examine emotional flexibility in the treatment of depression. *Administration and Policy in Mental Health and Mental Health Services Research*, 51:501–508.
- Cees Snoek, Marcel Worring, and Arnold W. M. Smeulders. 2005. Early versus late fusion in semantic video analysis. In *Proceedings of the 13th ACM International Conference on Multimedia, Singapore, November 6-11, 2005*, pages 399–402. ACM.
- Hao Sun, Jiaqing Liu, Shurong Chai, Zhaolin Qiu, Lanfen Lin, Xinyin Huang, and Yenwei Chen. 2021. Multi-modal adaptive fusion transformer network for the estimation of depression level. *Sensors*, 21(14).
- Hao Sun, Hongyi Wang, Jiaqing Liu, Yen-Wei Chen, and Lanfen Lin. 2022. Cubemlp: An mlp-based model for multimodal sentiment analysis and depression estimation. In *Proceedings of the 30th ACM International Conference on Multimedia, MM '22*, page 3722–3729, New York, NY, USA. Association for Computing Machinery.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6558–6569. Association for Computational Linguistics.
- Floris Van Steijn, Gizem Sogancioglu, and Heysem Kaya. 2022. Text-based interpretable depression severity modeling via symptom predictions. In *Proceedings of the 2022 International Conference on Multimodal Interaction, ICMI '22*, page 139–147, New York, NY, USA. Association for Computing Machinery.
- Jan Verwaeren, Willem Waegeman, and Bernard De Baets. 2012. Learning partial ordinal class memberships with kernel-based proportional odds models. *Computational Statistics & Data Analysis*, 56(4):928–942.
- Akkapon Wongkoblaph, Miguel A Vadillo, and Vasa Curcin. 2017. Researching mental health disorders in the era of social media: Systematic review. *J Med Internet Res*, 19(6):e228.
- Shi Yin, Cong Liang, Heyan Ding, and Shangfei Wang. 2019. A multi-modal hierarchical recurrent neural network for depression detection. In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop, AVEC '19*, page 65–71, New York, NY, USA. Association for Computing Machinery.
- Jeewoo Yoon, Chaewon Kang, Seungbae Kim, and Jinyoung Han. 2022. D-vlog: Multimodal vlog dataset for depression detection. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 12226–12234. AAAI Press.
- Chengbo Yuan, Xuxu Liu, Qianhui Xu, Yongqian Li, Yong Luo, and Xin Zhou. 2024. Depression diagnosis and analysis via multimodal multi-order factor fusion. In *Artificial Neural Networks and Machine Learning – ICANN 2024*, pages 56–70, Cham. Springer Nature Switzerland.
- Ziping Zhao and Keru Wang. 2022. Unaligned multimodal sequences for depression assessment from speech. In *44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society, EMBC 2022, Glasgow, Scotland, United Kingdom, July 11-15, 2022*, pages 3409–3413. IEEE.
- Bochao Zou, Jiali Han, Yingxue Wang, Rui Liu, Shenghui Zhao, Lei Feng, Xiangwen Lyu, and Huimin Ma. 2023. Semi-structural interview-based chinese multimodal depression corpus towards automatic preliminary screening of depressive disorders. *IEEE Trans. Affect. Comput.*, 14(4):2823–2838.

A PHQ-8 Questionnaire

In this section, we provide more details regarding the PHQ-8 questionnaire. The PHQ-8 questionnaire consists of the following questions:

- **Question 1:** Little interest or pleasure in doing things.
- **Question 2:** Feeling down, depressed, irritable or hopeless.
- **Question 3:** Trouble falling or staying asleep, or sleeping too much.
- **Question 4:** Feeling tired or having little energy.
- **Question 5:** Poor appetite or overeating.
- **Question 6:** Feeling bad about yourself – or that you are a failure or have let yourself or your family down.
- **Question 7:** Trouble concentrating on things, such as school work, reading or watching television.
- **Question 8:** Moving or speaking so slowly that other people could have noticed? Or the opposite – being so fidgety or restless that you have been moving around a lot more than usual.

These questions are scored from 0 to 3 based on how frequently the patients encounter them in the last two weeks. The scoring is based on the following:

- **Score 0:** Not at all
- **Score 1:** Several Days
- **Score 2:** More than Half Days
- **Score 3:** Nearly Everyday

The total score from all the questions is used to determine the depression severity of a patient. A higher score denotes higher depression severity.

B Evaluation Metrics

We use three different evaluation metrics for evaluation in this paper, which are elaborated below:

- **Concordance Correlation Coefficient (CCC):** CCC is a correlation based metric. It varies from -1 to 1 . A CCC of -1 between predicted values and actual values means the two variables are opposite. A CCC of 1 means they are identical, and 0 means they are not correlated. Thus, a higher CCC is desirable. CCC is defined as follows:

$$\rho_c = \frac{2\rho\sigma_x\sigma_y}{\sigma_x^2 + \sigma_y^2 + (\mu_x - \mu_y)^2} \quad (6)$$

where ρ is the pearson correlation between variables x and y . σ_x and σ_y are the standard deviations of variables x and y . μ_x and μ_y are the means of variables x and y .

- **Root Mean Squared Error (RMSE):** RMSE is a standard metric used in regression problems. It varies from 0 to ∞ . A lower RMSE is desirable. It is defined as follows:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N \|y(i) - x(i)\|^2} \quad (7)$$

- **Mean Absolute Error (MAE):** Mean Absolute error is another standard metric used in the evaluation of regression problems. It varies from 0 to ∞ . It is defined as follows:

$$MAE = \frac{1}{N} \sum_{i=1}^N \|y(i) - x(i)\| \quad (8)$$

C OpenSmile Low-level Features

Here, we provide more details about the low-level features extracted by OpenSmile, which are used in our experiments.

C.1 Frequency related parameters

Pitch: Logarithmic F0 on a semitone frequency scale, starting at 27.5 Hz. A semitone is the smallest music interval and is considered the most dissonant when sounded harmonically.

Jitter: Jitter measures the cycle-to-cycle variations of the fundamental frequency. It is a measure of frequency variability compared to the person's fundamental frequency.

Formant 1, 2, and 3 frequency: Centre frequency of first, second, and third formant. Formants are distinctive frequency components of the acoustic signal produced by speech. They are used to identify vowels.

Formant 1 Bandwidth: Bandwidth of the first formant. The formant bandwidth affects the identification of vowels in competition with other vowels.

C.2 Energy/Amplitude related parameters

Shimmer: Shimmer measures the cycle-to-cycle variations of fundamental amplitude. The shimmer changes with the reduction of glottal resistance and mass lesions on the vocal cords and is correlated with the presence of noise emission and breathiness.

Loudness: an estimate of perceived signal intensity from an auditory spectrum.

Harmonics-to-noise ratio: Relation of energy in harmonic components to energy in noise-like components. HNR quantifies the relative amount of additive noise.

C.3 Spectral Parameters

Alpha Ratio: Ratio of the summed energy from 50-1000 Hz and 1-5 kHz.

Hammarberg Index: Ratio of the strongest energy peak in the 0-2 kHz region to the strongest peak in the 2-5 kHz region.

Spectral Slope 0-500 Hz and 500-1500 Hz: Linear regression slope of the logarithmic power spectrum within the two given bands.

Formant 1, 2, and 3 relative energy: The ratio of the energy of the spectral harmonic peak at the first, second, and third formant’s centre frequency to the energy of the spectral peak at F0.

MFCC 1-4: Mel-Frequency Cepstral Coefficients 1-4.

Spectral flux: Difference of the spectra of two consecutive frames.

D *ImbOLL* and OLL parameters

Here, we show the experiments conducted to determine the optimal values for α and β for our *ImbOLL* function presented in equation 4. We experiment with $\alpha \in \{1, 1.5, 2\}$ and $\beta \in \{0.5, 1\}$. The results are shown in Table 4.

From Table 4, we see that $\alpha = 1$ gives the best results. Both $\beta = 0.5$ and $\beta = 1$ give similar results on mean performance. However, $\beta = 0.5$ gives a lower standard deviation. Thus, we choose the value of $\alpha = 1$ and $\beta = 0.5$ for training *QuestMF*.

We also experiment with the hyperparameters of the OLL function presented in equation 1. We experiment with $\alpha \in \{1, 1.5, 2\}$, and the results are presented in Table 5. From the table, we see

α	β	Validation CCC(\uparrow)
1	0.5	0.654 \pm 0.014
	1	0.653 \pm 0.022
1.5	0.5	0.639 \pm 0.026
	1	0.516 \pm 0.024
2	0.5	0.610 \pm 0.038
	1	0.422 \pm 0.024

Table 4: *ImbOLL* experiments

α	Validation CCC(\uparrow)
1	0.659 \pm 0.024
1.5	0.655 \pm 0.012
2	0.645 \pm 0.014

Table 5: OLL experiments

that $\alpha = 1$ and $\alpha = 1.5$ give good results. While $\alpha = 1$ gives slightly better mean performance, $\alpha = 1.5$ gives a lower standard deviation. So, we train *QuestMF* with both $\alpha = 1$ and $\alpha = 1.5$. The results are presented in Appendix F. From that, we observe that OLL with $\alpha = 1$ gives the best result on the three-modality fused model. Thus, only this result is presented in Section 6.

E Hyperparameter Details and Training Setup

This section presents the hyperparameters used in the single modality and fused two-modality and three-modality models. We use a maximum of 120 participant dialogue turns for all our experiments.

Single Modality Encoder Models: For the single modality models, we first experiment to find the best Bi-LSTM output dimensions. We experiment with hidden dimensions of $d \in \{30, 50, 100\}$. Due to the computational expenses, we only experimented with the text encoder model and extended the same output dimensions to audio and video encoder models. The results of this experiment are presented in Table 6 for the *QuestMF* framework trained with the *ImbOLL* function. The results of *QuestMF* framework trained with OLL with $\alpha = 1$ are presented in Table 7 and with OLL with $\alpha = 1.5$ are shown in Table 8. The results of *QuestMF* framework trained with the MSE loss function are presented in Table 9 and results of the *Total* framework are presented in Table 10. The tables show that a Bi-LSTM hidden dimension of 50 works best for all frameworks. For the multi-

Output Dimension	Validation CCC(\uparrow)
30	0.639 ± 0.027
50	0.654 ± 0.014
100	0.647 ± 0.004

Table 6: Results of experiments with LSTM output dimension for *QuestMF* Framework trained with *ImbOLL* function

Output Dimension	Validation CCC(\uparrow)
30	0.622 ± 0.047
50	0.659 ± 0.024
100	0.639 ± 0.038

Table 7: Results of experiments with LSTM output dimension for *QuestMF* Framework trained with OLL function with $\alpha = 1$

head attention layer, we use 4 attention heads and a dropout of 0.5 for the text encoder model for all frameworks. For the audio and video encoder models, we use two multihead attention layers with 4 attention heads and a dropout of 0.2 in all frameworks. The MLP in all single modality encoder models and frameworks consists of two linear layers with a hidden dimension of 256, and the ReLU activation function connects the linear layers. A dropout of 0.2 is applied before each linear layer.

For the training of the single modality encoders, we use a learning rate of 5×10^{-4} with AdamW optimiser and a batch size of 10 for all modalities. During training, we freeze the turn encoders and only train the Bi-LSTM layer, attention layer, and MLP. Since the text models fit faster, we train them for 20 epochs. Meanwhile, we train the audio and video models for 50 epochs. We select the model checkpoint with the lowest validation loss for further modality fusion training. To evaluate the models on the depression severity score prediction task, we select the checkpoint with the best validation CCC.

Two-Modality Fused Models: We follow the architecture shown in Figure 3 for the fusion of two modalities. We use 4 heads and a dropout of 0.8 for multihead cross-attention and self-attention layers. We use the very high dropout to reduce overfitting due to the small size of the training dataset. The MLP consists of two linear layers with a hidden dimension of 256. The linear layers are connected through the ReLU activation function. We apply a dropout of 0.8 before the first linear layer and a dropout of 0.5 before the last linear layer. We use a

Output Dimension	Validation CCC(\uparrow)
30	0.609 ± 0.028
50	0.655 ± 0.012
100	0.616 ± 0.015

Table 8: Results of experiments with LSTM output dimension for *QuestMF* Framework trained with OLL function with $\alpha = 1.5$

Output Dimension	Validation CCC(\uparrow)
30	0.554 ± 0.022
50	0.632 ± 0.024
100	0.602 ± 0.008

Table 9: Results of experiments with LSTM output dimension for *QuestMF* Framework trained with MSE loss function

Output Dimension	Validation CCC(\uparrow)
30	0.588 ± 0.015
50	0.614 ± 0.010
100	0.610 ± 0.012

Table 10: Results of experiments with LSTM output dimension for *Total* Framework

smaller dropout before the last linear layer to avoid underfitting.

For the training of a two-modality fusion encoder, we use a learning rate of 5×10^{-4} with AdamW optimiser and a batch size of 10. If the text modality is involved in the two-modality fusion model, we freeze the weights from the text encoder model while we train the weights in the audio or video encoder models during the fusion. This is because the text model fits the data quickly, so training the weights of the other model with the frozen text model helps information alignment across the modalities and improves their encodings. Another reason is that training the parameters of all single modality encoder models with the small training set would increase the chances of overfitting. In addition to this, we train the cross-attention layers and self-attention layers used for fusion and the MLP. We train the models for 20 epochs. We select the model checkpoint with the best validation CCC.

Three-Modality Fused Models: We follow the architecture shown in Figure 4 for the fusion of three modalities. We use 4 heads and a dropout of 0.8 in multihead cross-attention and self-attention layers. We use an MLP of two linear layers with a hidden dimension of 256. The linear layers are

connected through the ReLU activation function. We apply a dropout of 0.8 before the first linear layer and a dropout of 0.5 before the last linear layer.

For the training of the three-modality fused model, we use a learning rate of 5×10^{-4} with AdamW optimiser and a batch size of 10. While training three-modality fused models, we freeze the weights from the text model and train the weights of the audio and video models along with the cross-attention and self-attention layers used for fusion, and the MLP. We train the fusion for 20 epochs. We select the model checkpoint with the best validation CCC.

F Ablation Details

Here, we show more detailed results of our ablation study to observe the performance of single modality models, two-modality fused models, and the three-modality fused model with the frameworks. Here, we show the RMSE and MAE along with CCC results. We also show the standard deviation along with the mean for the three different seed runs. For *QuestMF* (OLL), we have two different frameworks here:

QuestMF (OLL-1): We train the *QuestMF* framework with OLL function with $\alpha = 1$.

QuestMF (OLL-1.5): We train the *QuestMF* framework with OLL function with $\alpha = 1.5$.

We present the results in Table 11.

Modalities	Framework	CCC(\uparrow)	RMSE(\downarrow)	MAE(\downarrow)
Text	<i>Total</i>	0.591 ± 0.031	5.51 ± 0.38	4.37 ± 0.30
	<i>QuestMF</i> (MSE)	0.593 ± 0.020	5.52 ± 0.09	4.33 ± 0.06
	<i>QuestMF</i> (OLL-1)	0.592 ± 0.025	5.77 ± 0.31	4.53 ± 0.22
	<i>QuestMF</i> (OLL-1.5)	0.616 ± 0.019	5.22 ± 0.07	4.02 ± 0.07
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.615 ± 0.031	5.71 ± 0.25	4.36 ± 0.26
Audio	<i>Total</i>	0.212 ± 0.017	6.41 ± 0.15	5.27 ± 0.09
	<i>QuestMF</i> (MSE)	0.239 ± 0.012	6.35 ± 0.09	5.14 ± 0.03
	<i>QuestMF</i> (OLL-1)	0.264 ± 0.008	6.96 ± 0.46	5.42 ± 0.20
	<i>QuestMF</i> (OLL-1.5)	0.256 ± 0.023	6.90 ± 0.21	5.36 ± 0.14
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.273 ± 0.021	6.67 ± 0.11	5.32 ± 0.05
Video	<i>Total</i>	-0.067 ± 0.009	8.21 ± 0.07	6.63 ± 0.04
	<i>QuestMF</i> (MSE)	-0.075 ± 0.007	7.97 ± 0.05	6.46 ± 0.04
	<i>QuestMF</i> (OLL-1)	-0.041 ± 0.015	7.91 ± 0.20	6.44 ± 0.18
	<i>QuestMF</i> (OLL-1.5)	-0.044 ± 0.026	7.79 ± 0.19	6.33 ± 0.16
	<i>QuestMF</i> (<i>ImbOLL</i>)	-0.052 ± 0.028	7.89 ± 0.12	6.44 ± 0.11
Text + Audio	<i>Total</i>	0.607 ± 0.020	5.27 ± 0.28	4.12 ± 0.23
	<i>QuestMF</i> (MSE)	0.618 ± 0.017	5.61 ± 0.31	4.42 ± 0.18
	<i>QuestMF</i> (OLL-1)	0.628 ± 0.013	5.44 ± 0.04	4.17 ± 0.04
	<i>QuestMF</i> (OLL-1.5)	0.622 ± 0.004	5.53 ± 0.07	4.27 ± 0.06
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.643 ± 0.024	5.48 ± 0.27	4.21 ± 0.14
Text + Video	<i>Total</i>	0.610 ± 0.008	5.13 ± 0.19	4.02 ± 0.17
	<i>QuestMF</i> (MSE)	0.627 ± 0.021	5.19 ± 0.16	3.99 ± 0.05
	<i>QuestMF</i> (OLL-1)	0.628 ± 0.006	5.34 ± 0.07	4.04 ± 0.20
	<i>QuestMF</i> (OLL-1.5)	0.630 ± 0.039	5.23 ± 0.31	4.05 ± 0.30
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.659 ± 0.018	5.22 ± 0.09	3.92 ± 0.08
Audio + Video	<i>Total</i>	0.058 ± 0.030	7.35 ± 0.10	5.86 ± 0.11
	<i>QuestMF</i> (MSE)	0.070 ± 0.011	7.35 ± 0.21	5.81 ± 0.16
	<i>QuestMF</i> (OLL-1)	0.139 ± 0.016	7.07 ± 0.11	5.69 ± 0.08
	<i>QuestMF</i> (OLL-1.5)	0.080 ± 0.047	7.42 ± 0.19	6.02 ± 0.20
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.159 ± 0.039	7.03 ± 0.18	5.67 ± 0.25
Text + Audio + Video	<i>Total</i>	0.602 ± 0.015	5.10 ± 0.08	3.99 ± 0.05
	<i>QuestMF</i> (MSE)	0.602 ± 0.025	5.36 ± 0.24	4.21 ± 0.12
	<i>QuestMF</i> (OLL-1)	0.640 ± 0.018	5.14 ± 0.05	3.88 ± 0.01
	<i>QuestMF</i> (OLL-1.5)	0.599 ± 0.023	5.30 ± 0.22	4.06 ± 0.25
	<i>QuestMF</i> (<i>ImbOLL</i>)	0.662 ± 0.022	5.25 ± 0.08	3.95 ± 0.13

Table 11: Ablation results for using different combinations of modalities with different frameworks. The CCC, RMSE and MAE scores presented are the mean and standard deviation over 3 different seed runs.