

# FACTORIZED MULTIMODAL TRANSFORMER FOR MULTIMODAL SEQUENTIAL LEARNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The complex world around us is inherently multimodal and sequential (continuous). Information is scattered across different modalities and requires multiple continuous sensors to be captured. As machine learning leaps towards better generalization to real world, multimodal sequential learning becomes a fundamental research area. Arguably, modeling arbitrarily distributed spatio-temporal dynamics within and across modalities is the biggest challenge in this research area. In this paper, we present a new transformer model, called the Factorized Multimodal Transformer (FMT) for multimodal sequential learning. FMT inherently models the intramodal and intermodal (involving two or more modalities) dynamics within its multimodal input in a factorized manner. The proposed factorization allows for increasing the number of self-attentions to better model the multimodal phenomena at hand; without encountering difficulties during training (e.g. overfitting) even on relatively low-resource setups. All the attention mechanisms within FMT have a full time-domain receptive field which allows them to asynchronously capture long-range multimodal dynamics. In our experiments we focus on datasets that contain the three commonly studied modalities of language, vision and acoustic. We perform a wide range of experiments, spanning across 3 well-studied datasets and 21 distinct labels. FMT shows superior performance over previously proposed models, setting new state of the art in the studied datasets.

## 1 INTRODUCTION

In many naturally occurring scenarios, our perception of the world is multimodal. For example, consider multimodal language (face-to-face communication), where modalities of language, vision and acoustic are seamlessly used together for communicative intent (Kottur et al., 2019). Such scenarios are widespread in everyday life, where continuous sensory perceptions form multimodal sequential data. Each modality within multimodal data exhibits exclusive *intramodal* dynamics, and presents a unique source of information. Modalities are not fully independent of each other. Relations across two (bimodal) or more (trimodal, ...) of them form *intermodal* dynamics; often asynchronous spatio-temporal dynamics which bind modalities together (Zadeh et al., 2017).

Learning from multimodal sequential data has been an active, yet challenging research area within the field of machine learning (Baltrušaitis et al., 2018). Various approaches relying on graphical models or RNNs have been proposed for multimodal sequential learning. Transformer models are a new class of neural models that rely on a carefully designed non-recurrent architecture for sequential modeling (Vaswani et al., 2017). Their superior performance is attributed to a self-attention mechanism, which is uniquely capable of highlighting related information across a sequence. This self-attention is a particularly appealing mechanism for multimodal sequential learning, as it can be modified into a strong neural component for finding relations between different modalities (the cornerstone of this paper). In practice, numerous such relations may simultaneously exist within multimodal data, which would require increasing the number of attention units (i.e. heads). Increasing the number of attentions in an efficient and semantically meaningful way inside a transformer model, can boost the performance in modeling multimodal sequential data.

In this paper, we present a new transformer model for multimodal sequential learning, called Factorized Multimodal Transformer (FMT). FMT is capable of modeling asynchronous intramodal and intermodal dynamics in an efficient manner, within one single transformer network. It does so by

specifically accounting for possible sets of interactions between modalities (i.e. factorizing based on combinations) in a Factorized Multimodal Self-attention (FMS) unit. We evaluate the performance of FMT on multimodal language: a challenging type of multimodal data which exhibits idiosyncratic and asynchronous spatio-temporal relations across language, vision and acoustic modalities. FMT is compared to previously proposed approaches for multimodal sequential learning over multimodal sentiment analysis (CMU-MOSI) (Zadeh et al., 2016), multimodal emotion recognition (IEMOCAP) (Busso et al., 2008), and multimodal personality traits recognition (POM) (Park et al., 2014).

## 2 RELATED WORKS

The related works to studies in this paper fall into two main areas.

### 2.1 MULTIMODAL SEQUENTIAL LEARNING

Modeling multimodal sequential data is among the core research areas within the field of machine learning. In this area, previous work can be classified into two main categories.

The first category of models, and arguably the simplest, are models that use early or late fusion. Early fusion uses feature concatenation of all modalities into a single modality. Subsequently, the multimodal sequential learning task is treated as a unimodal one and tackled using unimodal sequential models such as Hidden Markov Models (HMMs) (Baum & Petrie, 1966), Hidden Conditional Random Fields (HCRFs) (Quattoni et al., 2007; Morency et al., 2007), and RNNs (e.g. LSTMs Hochreiter & Schmidhuber (1997)). While such models are often successful for real unimodal data (i.e. not feature concatenated multimodal data), they lack the necessary components to deal with multimodal data often causes suboptimal performance (Xu et al., 2013). Contrary to early fusion which concatenates modalities at input level, late fusion models have relied on learning ensembles of weak classifiers from different modalities (Snoek et al., 2005; Vielzeuf et al., 2017; Nojavanasghari et al., 2016). Hybrid methods have also been used to combine early and late fusion together (Wu et al., 2019; Nguyen & Okatani, 2018; Lazaridou et al., 2015; Hu et al., 2017).

The second category of models comprise of models specifically designed for multimodal data. Multimodal variations of graphical models have been proposed, including Multi-view HCRFs where the potentials of the HCRF are changed to facilitate multiple modalities Song et al. (2012; 2013). Multimodal models based on LSTMs include Multi-view LSTMs Rajagopalan et al. (2016), Memory Fusion Network (Zadeh et al., 2018a) with its recurrent and graph variants (Liang et al., 2018; Zadeh et al., 2018c), as well as Multi-attention Recurrent Networks (Zadeh et al., 2018b). Studies have also proposed generic fusion techniques that can be used in various models including Tensor Fusion (Zadeh et al., 2017) and its approximate variants (Liang et al.; Liu et al., 2018), as well as Compact Bilinear Pooling (Gao et al., 2015; Fukui et al., 2016; Kim et al., 2016).

Many of these models, from both first and second categories, are used as baselines in this paper.

### 2.2 TRANSFORMER MODEL

Transformer is a non-recurrent neural architecture designed for modeling sequential data (Vaswani et al., 2017). It has shown superior performance across multiple NLP tasks when compared to RNN-based or convolutional architectures (Devlin et al., 2018; Vaswani et al., 2017). This superior performance of Transformer model is largely credited to a self-attention; a neural component that allows for efficiently extracting both short and long-range dependencies within its input sequence space. Transformer models have been successfully applied to various areas within machine learning including NLP and computer vision (Yang et al., 2019; Parmar et al., 2018; Alsentzer et al., 2019). Extending transformer to multimodal domains, specially for structured multimodal sequences is relatively understudied; with the previous works mainly focusing on using transformer models for modality alignment using cross-modal links between single transformers for each modality (Tsai et al., 2019).

### 3 FACTORIZED MULTIMODAL TRANSFORMER (FMT) MODEL

In this section, we outline the proposed Factorized Multimodal Transformer<sup>1</sup> (FMT). Figure 1 shows the overall structure of the FMT model. The input first goes through an embedding layer, followed by multiple Multimodal Transformer Layers (MTL). Each MTL consists of multiple Factorized Multimodal Self-attentions (FMS). FMS explicitly accounts for intramodal and intermodal factors within its multimodal input.  $S1$  and  $S2$  are two summarization networks. They are necessary components of FMT which allow for increasing the number of attentions efficiently, without overparameterization of the FMT.

#### 3.1 INPUT EMBEDDING

Consider a multimodal sequential dataset with constituent modalities of language, vision and acoustic. The modalities are denoted as  $\{L, V, A\}$  from hereon for abbreviation. After resampling using a reference clock, modalities can follow the same frequency (Chen et al., 2017). Essentially, this resampling is often based on word timestamps (i.e. word alignment). Subsequently, the dataset can be denoted as:

$$D = \left\{ x_i = \left[ x_{(t,i)} = \langle l_{(t,i)}, v_{(t,i)}, a_{(t,i)} \rangle \right]_{t=1}^{\mathcal{T}_i}, y_i \right\}_{i=1}^N$$

$x_i \in \mathbb{R}^{\mathcal{T}_i \times d_x}$ ,  $y_i \in \mathbb{R}^{d_y}$  are the inputs and labels.  $x_{(t,i)} = \langle l_{(t,i)}, v_{(t,i)}, a_{(t,i)} \rangle$  is a triplet of language, visual and audio inputs for timestamp  $t$  in  $i$ -th datapoint.  $N$  is the total number of samples within the dataset, and  $\mathcal{T}_i$  the total number of timestamps within  $i$ -th datapoint. Zero paddings (on the left) can be used to unify the length of all sequences to a desired fixed length  $\mathcal{T}$ .  $d_x = d_L + d_V + d_A$  denotes the dimensionality of input at each timestep, which in turn is equal to the sum of dimensionality of each modality.  $d_y$  denotes the dimensionality of the associated labels of a sequence.

At the first step within the FMT model, each modality is passed to a unimodal embedding layer with the operation  $\mathcal{E}_{M \in \{L, V, A\}}(\cdot); \mathbb{R}^{d_{M \in \{L, V, A\}}} \mapsto \mathbb{R}^{e_{M \in \{L, V, A\}}}$ . In turn,  $\mathcal{E}_M$  takes as input  $m_{(t,i)}$ ;  $m \in \{l, v, a\}$ . Positional embeddings are also added to the input at this stage. The output of the embeddings collectively form  $\hat{x}^0 = \langle \hat{l}^0, \hat{v}^0, \hat{a}^0 \rangle$ . We denote the dimensionality of this output as  $e_x = e_L + e_V + e_A$ .

#### 3.2 MULTIMODAL TRANSFORMER LAYER (MTL)

After the initial embedding, FMT now consists of a stack of Multimodal Transformer Layers (MTL). MTL 1) captures factorized dynamics within multimodal data in parallel, and 2) aligns the time-asynchronous information both within and across modalities. Both of these are achieved using multiple Factorized Multimodal Self-attentions (FMS), each of which has multiple specialized self-attentions inside. The high dimensionality of the intermediate attention outputs within MTL and FMS is controlled using two distinct summarization networks. The continuation of this section provides detailed explanation of the inner-operations of MTL.

Let  $\hat{x}_i^k = \langle \hat{l}_{(i,i)}^k, \hat{v}_{(i,i)}^k, \hat{a}_{(i,i)}^k \rangle$  denote the input to the  $k$ -th MTL. We assume a total of  $K$  MTLs in a FMT (indexed  $0 \dots K-1$ ), with  $k=0$  being the output of the embedding layer (input to  $k=0$  MTL). The input of MTL, immediately goes through one/multiple<sup>2</sup> Factorized

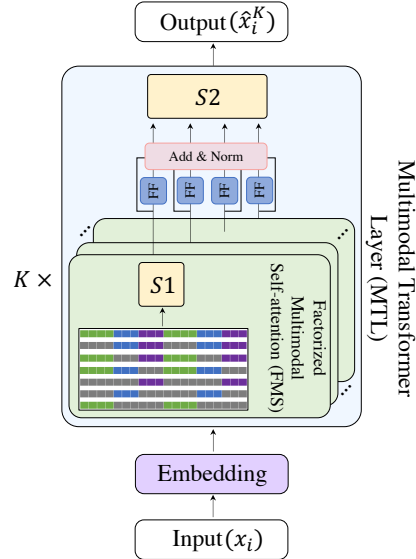


Figure 1: Overview of the proposed Factorized Multimodal Transformer (FMT) model.

<sup>1</sup>Code: [github.com/removed-for-blind-review](https://github.com/removed-for-blind-review), Public Data: <https://github.com/A2Zadeh/CMU-MultimodalSDK>

<sup>2</sup>Multiple FMS have the same time-domain receptive field, which is equal to the length of the input. This is contrary to the implementations of the transformer model that split the sequence based on number of attention heads.




Figure 2: Best viewed in color. Overview of a single Factorized Multimodal Self-attention (FMS) in  $k$ -th MTL. SA is self-attention (Vaswani et al., 2017) with full time-domain receptive field ( $0 \dots \mathcal{T}$ ). The grayed areas are for demonstration purposes, and not a part of the implementation.

Multimodal Self-attentions (FMS). The operations inside a single Factorized Multimodal Self-attention is demonstrated in Figure 2. For 3 modalities<sup>3</sup>, there exist 7 distinct attentions inside a single FMS unit. Each attention has a unique receptive field with respect to modalities  $f \in F = \{L, V, A, LV, LA, VA, LVA\}$ ; essentially denoting the modalities visible to the attention. Using this factorization, FMS explicitly accounts for possible unimodal, bimodal and trimodal interactions existing within the multimodal input space. All attentions within a FMS extend to the length of the sequence, and therefore can extract asynchronous relations within and across modalities. For  $f \in F$ , each attention within a single FMS unit is controlled by the Key  $K^f$ , Query  $Q^f$ , and Value  $V^f$  all with dimensionality  $\mathbb{R}^{\mathcal{T} \times \mathcal{T}}$ ; parameterized respectively using affine maps  $W_{K^f}$ ,  $W_{Q^f}$ , and  $W_{V^f}$ . After the attention is applied using Key, Query and Value operations (Vaswani et al., 2017), the output of each of the attentions goes through a residual addition with its perceived input (input in the attention receptive field), followed by a normalization.

The output of the FMS contains the aligned and extracted information from the unimodal, bimodal and trimodal factors. This output is high-dimensional; essentially  $\mathbb{R}^{4 \times \mathcal{T} \times e_x}$  (each dimension within input of shape  $\mathcal{T} \times e_x$  is present in 4 factors). Our goal is to reduce this high-dimensional data using a mapping from  $\mathbb{R}^{4 \times \mathcal{T} \times e_x} \mapsto \mathbb{R}^{\mathcal{T} \times e_x}$ . Without overparameterizing the FMS, in practice, we observed this mapping can be efficiently done using a simple 1D convolutional network  $\mathcal{S}_{1_{M \in \{L, V, A\}}}(\cdot); \mathbb{R}^4 \mapsto \mathbb{R}$ . Internally,  $\mathcal{S}_{1_L}(\cdot)$  maps its input to multiple layers of higher dimensions and subsequently to  $\mathbb{R}$ . Using language as an example,  $\mathcal{S}_{1_L}$  moves across language modality dimensions  $e_L$  for  $t = 1 \dots \mathcal{T}$  and summarizes the information across all the factors. The output of this summarization applied on all modality dimensions and timesteps, is the output of FMS, which has the dimensionality  $\mathbb{R}^{\mathcal{T} \times e_x}$ .

In practice, there can be various possible unimodal, bimodal or trimodal interactions within a multimodal input. For example, consider multiple sets of important interactions between L and V (e.g. smile + positive word, as well as eyebrows up + excited phrase), all of which need to be highlighted and extracted. A single FMS may not be able to highlight all these interactions without diluting its intrinsic attentions. Multiple FMS can be used inside a MTL to efficiently extract diverse multimodal interactions existing in the input data<sup>4</sup>. Consider a total of  $U$  FMS units inside a MTL. The output of each FMS goes through a feedforward network (for each timestamp  $t$  of the FMS output).

<sup>3</sup>Remarks on more than 3 modalities is in Appendix A.4.

<sup>4</sup>We study the impact of number of FMS units inside MTL in Section 5.

<i>Model \ Metric</i>	<i>BA</i>	<i>F1</i>	<i>MAE</i>	<i>Corr</i>
MV-LSTM (Rajagopalan et al., 2016)	73.9/-	74.0/-	1.019	0.601
TFN (Zadeh et al., 2017)	73.9/-	73.4/-	1.040	0.633
MARN (Zadeh et al., 2018b)	77.1/-	77.0/-	0.968	0.625
MFN (Zadeh et al., 2018a)	77.4/-	77.3/-	0.965	0.632
RMFN (Liang et al., 2018)	78.4/-	78.0/-	0.922	0.681
RAVEN (Wang et al., 2018)	78.0/-	-/-	0.915	0.691
MulT (Tsai et al., 2019)	-/83.0	-/82.8	0.87	0.698
FMT (ours)	<b>81.5/83.5</b>	<b>81.4/83.5</b>	<b>0.837</b>	<b>0.744</b>

Table 1: FMT achieves superior performance over baseline models for CMU-MOSI dataset (multimodal sentiment analysis). We report BA (binary accuracy) and F1 (both higher is better), MAE (Mean-absolute Error, lower is better), and Corr (Pearson Correlation Coefficient, higher is better). For BA and F1, we report two numbers: the number on the left side of “/” is calculated based on approach taken by Zadeh et al. (2018b), and the right side is by Tsai et al. (2019).

The output of this feedforward network is residually added with its input, and subsequently normalized. The feedforward network is the same across all  $U$  FMS units and timestamps  $t$ . Subsequently, the dimensionality of the output of the normalizations collectively is  $\mathbb{R}^{U \times T \times e_x}$ . Similar to operations performed by  $\mathcal{S}1$ , a secondary summarization network  $\mathcal{S}2_{\mathcal{M} \in \{\text{L}, \text{V}, \text{A}\}}(\cdot); \mathbb{R}^U \mapsto \mathbb{R}$  can be used here.  $\mathcal{S}2$  is also a 1D convolutional network that moves across modality dimensions and different timesteps to map  $\mathbb{R}^{U \times T \times e_x}$  to  $\mathbb{R}^{T \times e_x}$ . The output of the secondary summarization network is the final output of MTL, and denoted as  $\hat{x}_i^{k+1}$ .

Let  $\hat{x}_i^K = \langle \hat{l}_{(\cdot, i)}^K, \hat{v}_{(\cdot, i)}^K, \hat{a}_{(\cdot, i)}^K \rangle$  be the output of last MTL in the stack. For supervision, we feed this input one timestamp at a time as input to a Gated Recurrent Unit (GRU) (Cho et al., 2014). The prediction is conditioned on output at timestamp  $t = T$  of the GRU, using an affine map to  $d_y$ .

## 4 EXPERIMENTAL METHODOLOGY

In this section, we discuss the experimental methodology including tasks, datasets, computational descriptors, and comparison baselines.

### 4.1 TASKS AND DATASETS

The following inherently multimodal tasks (and accompanied datasets) are studied in this paper. All the tasks are related to multimodal language: a complex and idiosyncratic sequential multimodal signal, where semantics are arbitrarily scattered across modalities (Holler & Levinson, 2019).

*Multimodal Sentiment Analysis:* The first benchmark in our experiments is multimodal sentiment analysis, where the goal is to identify a speaker’s sentiment based on the speaker’s display of verbal and nonverbal behaviors. We use the well-studied CMU-MOSI (CMU Multimodal Opinion Sentiment Intensity) dataset for this purpose (Zadeh et al., 2016). There are a total of 2199 data points (opinion utterances) within CMU-MOSI dataset. The dataset has real-valued sentiment intensity annotations in the range  $[-3, +3]$ . It is considered a challenging dataset due to speaker diversity (1 video per distinct speaker), topic variations and low-resource setup.

*Multimodal Emotion Recognition:* The second benchmark in our experiments is multimodal emotion recognition, where the goal is to identify a speaker’s emotions based on the speaker’s verbal and nonverbal behaviors. We use the well-studied IEMOCAP dataset (Busso et al., 2008). IEMOCAP consists of 151 sessions of recorded dialogues, of which there are 2 speaker’s per session for a total of 302 videos across the dataset. We perform experiments for discrete emotions (Ekman, 1992) of Happy, Sad, Angry and Neutral (no emotions) - similar to previous works (Tsai et al., 2019; Wang et al., 2018).

<i>Model \ Emotion</i>	<i>Happy</i>		<i>Sad</i>		<i>Angry</i>		<i>Neutral</i>	
	<i>BA</i>	<i>F1</i>	<i>BA</i>	<i>F1</i>	<i>BA</i>	<i>F1</i>	<i>BA</i>	<i>F1</i>
MV-LSTM (Rajagopalan et al., 2016)	85.9	81.3	80.4	74.0	85.1	84.3	67.0	66.7
MARN (Zadeh et al., 2018a)	86.7	83.6	82.0	81.2	84.6	84.2	66.8	65.9
MFN (Zadeh et al., 2018a)	86.5	84.0	83.5	82.1	85.0	83.7	69.6	69.2
RMFN (Liang et al., 2018)	87.5	85.8	82.9	85.1	84.6	84.2	69.5	69.1
RAVEN (Wang et al., 2018)	87.3	85.8	83.4	83.1	87.3	86.7	69.7	69.3
MuT (Tsai et al., 2019)	<b>90.7</b>	<b>88.6</b>	86.7	86.0	87.4	87.0	72.4	70.7
FMT	88.8	87.2	<b>88.0</b>	<b>87.7</b>	<b>89.7</b>	<b>89.5</b>	<b>74.0</b>	<b>73.8</b>

Table 2: FMT achieves superior performance over baseline models (with the exception of Happy emotion) for discrete emotions in IEMOCAP (multimodal emotion recognition). We report BA (binary accuracy) and F1 (both higher is better).

<i>Model \ Trait</i>	<i>Con</i>	<i>Pas</i>	<i>Voi</i>	<i>Dom</i>	<i>Cre</i>	<i>Viv</i>	<i>Exp</i>	<i>Ent</i>
	<i>MA7</i>	<i>MA7</i>	<i>MA7</i>	<i>MA7</i>	<i>MA7</i>	<i>MA7</i>	<i>MA7</i>	<i>MA7</i>
MV-LSTM (Rajagopalan et al., 2016)	25.6	28.6	28.1	34.5	25.6	32.5	32.5	29.6
TFN (Zadeh et al., 2017)	24.1	31.0	31.5	34.5	24.6	25.6	27.6	29.1
MARN (Zadeh et al., 2018b)	29.1	33.0	-	-	31.5	-	-	-
MFN (Zadeh et al., 2018a)	34.5	35.5	37.4	41.9	34.5	36.9	36.0	37.9
RMFN (Liang et al., 2018)	37.4	38.4	37.4	-	37.4	38.9	38.9	-
MuT (Tsai et al., 2019)	34.5	34.5	36.5	38.9	37.4	36.9	37.9	39.4
FMT	<b>40.9</b>	<b>42.4</b>	<b>42.4</b>	<b>44.3</b>	<b>41.4</b>	<b>39.4</b>	<b>41.4</b>	<b>39.4</b>

<i>Model \ Trait</i>	<i>Res</i>	<i>Tru</i>	<i>Rel</i>	<i>Out</i>	<i>Tho</i>	<i>Ner</i>	<i>Per</i>	<i>Hum</i>
	<i>MA5</i>	<i>MA5</i>	<i>MA5</i>	<i>MA5</i>	<i>MA5</i>	<i>MA5</i>	<i>MA7</i>	<i>MA5</i>
MV-LSTM (Rajagopalan et al., 2016)	33.0	52.2	50.7	38.4	37.9	42.4	26.1	38.9
TFN (Zadeh et al., 2017)	30.5	38.9	35.5	37.4	33.0	42.4	27.6	33.0
MARN (Zadeh et al., 2018b)	36.9	-	52.2	-	-	47.3	31.0	44.8
MFN (Zadeh et al., 2018a)	38.4	57.1	53.2	46.8	47.3	47.8	34.0	47.3
RMFN (Liang et al., 2018)	39.4	-	53.7	-	48.3	48.3	35.0	46.8
MuT (Tsai et al., 2019)	41.4	60.6	54.2	43.3	49.3	46.3	33.5	43.3
FMT	<b>44.8</b>	<b>61.1</b>	<b>57.6</b>	<b>51.7</b>	<b>51.7</b>	<b>51.2</b>	<b>40.4</b>	<b>48.3</b>

Table 3: FMT achieves superior performance over baseline models in POM dataset (multimodal personality traits recognition). For label abbreviations please refer to Section 4.3. MA(5,7) denotes multi-class accuracy for (5,7)-class personality labels (higher is better).

*Multimodal Speaker Traits Recognition:* The third benchmark in our experiments is speaker trait recognition based on communicative behavior of a speaker. It is a particularly difficult task, with 16 different speaker traits in total. We study the POM dataset which contains 1,000 movie review videos (Park et al., 2014). Each video is annotated for various personality and speaker traits, specifically: Confident (Con), Passionate (Pas), Voice Pleasant (Voi), Dominant (Dom), Credible (Cre), Vivid (Viv), Expertise (Exp), Entertaining (Ent), Reserved (Res), Trusting (Tru), Relaxed (Rel), Outgoing (Out), Thorough (Tho), Nervous (Ner), Persuasive (Per) and Humorous (Hum). The short form of these speaker traits is indicated inside the parentheses and used for the rest of this paper.

#### 4.2 MULTIMODAL COMPUTATIONAL DESCRIPTORS

The following computational descriptors are used by FMT and baselines (all the baselines use the same descriptors in their original respective papers).

*Language*: P2FA forced alignment model (Yuan & Liberman, 2008) is used to align the text and audio at word level. From the forced alignment, the timing of words and sentences are extracted. Word-level alignment is used to unify the modality frequencies (Chen et al., 2017). GloVe embeddings (Pennington et al., 2014) are subsequently used for word representation.

*Visual*: For the visual modality, the Emotient FACET (iMotions, 2017) is used to extract a set of visual features including Facial Action Units (Ekman et al., 1980), visual indicators of emotions, and sparse facial landmarks.

*Acoustic*: COVAREP (Degottex et al., 2014) is used to extract the following features: fundamental frequency, quasi open quotient (Kane & Gobl, 2013), normalized amplitude quotient, glottal source parameters (H1H2, Rd, Rd conf) (Drugman et al., 2012), Voiced/Unvoiced segmenting features (VUV) (Drugman & Alwan, 2011), maxima dispersion quotient (MDQ), the first 3 formants, parabolic spectral parameter (PSP), harmonic model and phase distortion mean (HMPDM 0-24) and deviations (HMPDD 0-12), spectral tilt/slope of wavelet responses (peak/slope), Mel Cepstral Coefficients (MCEP 0-24).

### 4.3 BASELINE MODELS AND PERFORMANCE MEASURES

The following strong baselines are compared to FMT: *MV-LSTM* (*Multi-view LSTM*, Rajagopalan et al. (2016)), *TFN* (*Tensor Fusion Network*, Zadeh et al. (2017)), *MARN* (*Multi-attention Recurrent Network*, Zadeh et al. (2018b)), *MFN* (*Memory Fusion Network*, Zadeh et al. (2018a)), *RAVEN* (*Recurrent Attended Variation Embedding Network*, Wang et al. (2018)), *MuT<sup>5</sup>* (*Multimodal Transformer for [Un]aligned Sequences*, Tsai et al. (2019)). There are fundamental distinctions between FMT and MuT, chief among them: 1) MuT consists of 6 transformers, 3 cross-modal transformers and 3 unimodal. Naturally this increases the overall model size substantially. FMT consists of only one transformer, with components to avoid overparameterization. 2) FMT sees interactions as undirected (unlike MuT which has  $L \rightarrow V$  and  $V \rightarrow L$ ), and therefore semantically combines two attentions in one. 3) MuT has no trimodal factors (which are important according to Section 5). 4) MuT has no direct unimodal path (e.g. only  $L$ ), as input to unimodal transformers are outputs of cross-modal transformers. 5) All FMT attentions have full time-domain receptive field, while MuT splits the input based on the heads.

In their original publication, all the models report<sup>6</sup> the performance over the datasets in Section 4.1, using the same descriptors discussed in Section 4.2. The models in this paper are compared using the following performance measures (depending on the dataset): (BA) denotes binary accuracy - higher is better, (MA5,MA7) are 5 and 7 multiclass accuracy - higher is better, (F1) denotes F1 score - higher is better, (MAE) denotes the Mean-Absolute Error - lower is better, (Corr) is Pearson Correlation Coefficient - higher is better. The hyperparameter space search for FMT (and baselines if retrained) is discussed in Appendix A.1.

## 5 RESULTS AND DISCUSSION

The results of sentiment analysis experiments on CMU-MOSI dataset are presented in Table 1. FMT achieves superior performance than the previously proposed models for multimodal sentiment analysis. We use two approaches for calculating BA and F1 based on negative vs. non-negative sentiment (Zadeh et al., 2018b) on the left side of /, and negative vs. positive (Tsai et al., 2019) on the right side. MAE and Corr are also reported. For multimodal emotion recognition, experiments on IEMOCAP are reported in Table 2. The performance of FMT is superior than other baselines for multimodal emotion recognition (with the exception of Happy emotion). The results of experiments for personality traits recognition on POM dataset are reported in Table 3. We report MA5 and MA7, depending on the label. FMT outperforms baselines across all personality traits.

We study the importance of the factorization in FMT. We first remove the unimodal, bimodal and trimodal attentions from the FMT model, resulting in 3 alternative implementations of FMT. Table 4

<sup>5</sup>We use the aligned variant of MuT model, which has shown better performance than unaligned version in the original paper.

<sup>6</sup>With the exception of MuT for POM dataset, which is not reported in original paper. It is trained in this paper using authors' provided github code with hyperparameter search in Appendix A.1.

<i>Model \ Metric</i>	<i>BA</i>	<i>F1</i>	<i>MAE</i>	<i>Corr</i>
FMT [UNI]	80.6/82.8	80.5/82.8	0.868	0.719
FMT [BI]	81.2/81.7	81.2/81.6	0.877	0.706
FMT [TRI]	80.2/81.6	80.2/81.5	0.874	0.705
FMT [L]	77.7/79.6	77.7/79.6	0.935	0.666
FMT [A]	62.5/62.7	62.6/73.2	1.338	0.306
FMT [V]	59.3/59.3	59.4/72.7	1.357	0.218
FMT [S]	80.3/82.0	80.3/81.9	0.860	0.734
FMT	<b>81.5/83.5</b>	<b>81.4/83.5</b>	<b>0.837</b>	<b>0.744</b>

Table 4: FMT ablation studies on CMU-MOSI dataset. UNI, BI, TRI denote removing all unimodal, bimodal and trimodal factors respectively. L, A, V denote using only language, audio, and visual factors respectively. S denotes the model with summarization networks replaced by simple addition. All factors, modalities, and components are needed for achieving best performance.

<i>Model \ Metric</i>	<i>BA</i>	<i>F1</i>	<i>MAE</i>	<i>Corr</i>
FMT [1]	80.0/82.0	79.4/81.1	0.864	0.712
FMT [2]	79.7/82.2	79.7/82.2	0.863	0.725
FMT [3]	79.2/80.9	79.1/80.8	0.905	0.698
FMT [4]	79.4/81.1	79.4/81.0	0.855	0.733
FMT [5]	81.5/82.6	81.5/82.5	0.886	0.711
FMT [6]	<b>81.5/83.5</b>	<b>81.4/83.5</b>	<b>0.837</b>	<b>0.744</b>

Table 5: Multimodal sentiment analysis experiments with different number of FMS units inside MTL. The number in the square bracket indicates the number of FMS. The total number of attention is 7 times the number in bracket.

demonstrates the results of this ablation experiment over CMU-MOSI dataset. Furthermore, we use only one modality as input for FMT, to understand the importance of each modality (all other factors removed). We also replace the summarization networks with simple vector addition operation. All factors, modalities, and summarization components are needed for achieving best performance.

We also perform experiments to understand the effect of number of FMT units within each MTL. Table 5 shows the performance trend for different number of FMT units. The model with 6 number of FMS (42 attentions in total) achieves the highest performance (6 is also the highest number we experimented with). Tsai et al. (2019) reports the best performance for CMU-MOSI dataset is achieved when using 40 attentions per cross-modal transformer (3 of each, therefore 120 attention, without counting the subsequent unimodal transformers). FMT uses fewer number of attentions than MulT, yet achieves better performance. We also experiment with number of heads for original transformer model (Vaswani et al., 2017) and compare to FMT (Appendix A.3).

## 6 CONCLUSION

In this paper, we presented the Factorized Multimodal Transformer (FMT) model for multimodal sequential learning. Using a Factorized Multimodal Self-attention (FMS) within each Multimodal Transformer Layer (MTL), FMT is able to model the intra-model and inter-modal dynamics within asynchronous multimodal sequences. We compared the performance of FMT to baselines approaches over 3 publicly available datasets for multimodal sentiment analysis (CMU-MOSI, 1 label), emotion recognition (IEMOCAP, 4 labels) and personality traits recognition (POM, 16 labels). Overall, FMT achieved superior performance than previously proposed models across the studied datasets.



## REFERENCES

- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. Publicly available clinical BERT embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pp. 72–78, Minneapolis, Minnesota, USA, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-1909. URL <https://www.aclweb.org/anthology/W19-1909>.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, 2018.
- Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette Chang, Sungbok Lee, and Shrikanth S. Narayanan. Iemocap: Interactive emotional dyadic motion capture database. *Journal of Language Resources and Evaluation*, 42(4):335–359, dec 2008. doi: 10.1007/s10579-008-9076-6.
- Minghai Chen, Sen Wang, Paul Pu Liang, Tadas Baltrušaitis, Amir Zadeh, and Louis-Philippe Morency. Multimodal sentiment analysis with word-level fusion and reinforcement learning. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pp. 163–171. ACM, 2017.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer. Covarepa collaborative voice analysis repository for speech technologies. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pp. 960–964. IEEE, 2014.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018. URL <http://arxiv.org/abs/1810.04805>.
- Thomas Drugman and Abeer Alwan. Joint robust voicing detection and pitch estimation based on residual harmonics. In *Interspeech*, pp. 1973–1976, 2011.
- Thomas Drugman, Mark Thomas, Jon Gudnason, Patrick Naylor, and Thierry Dutoit. Detection of glottal closure instants from speech signals: A quantitative review. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(3):994–1006, 2012.
- Paul Ekman. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200, 1992.
- Paul Ekman, Wallace V Friesen, and Sonia Ancoli. Facial signs of emotional experience. *Journal of personality and social psychology*, 39(6):1125, 1980.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*, 2016.
- Yang Gao, Oscar Beijbom, Ning Zhang, and Trevor Darrell. Compact bilinear pooling. *CoRR*, abs/1511.06062, 2015. URL <http://arxiv.org/abs/1511.06062>.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Judith Holler and Stephen C Levinson. Multimodal language processing in human communication. *Trends in Cognitive Sciences*, 2019.
- Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to reason: End-to-end module networks for visual question answering. *CoRR*, abs/1704.05526, 2017. URL <http://arxiv.org/abs/1704.05526>.
- iMotions. Facial expression analysis, 2017. URL [goo.gl/1rh1JN](http://goo.gl/1rh1JN).
- John Kane and Christer Gobl. Wavelet maxima dispersion for breathy to tense voice discrimination. *IEEE Transactions on Audio, Speech, and Language Processing*, 21(6):1170–1179, 2013.

- Jin-Hwa Kim, Kyoung Woon On, Woosang Lim, Jeonghee Kim, JungWoo Ha, and Byoung-Tak Zhang. Hadamard product for low-rank bilinear pooling. *CoRR*, abs/1610.04325, 2016. URL <http://arxiv.org/abs/1610.04325>.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- David G Kleinbaum, Lawrence L Kupper, Keith E Muller, and Azhar Nizam. *Applied regression analysis and other multivariable methods*, volume 601. Duxbury Press Belmont, CA, 1988.
- Satwik Kottur, José M. F. Moura, Devi Parikh, Dhruv Batra, and Marcus Rohrbach. CLEVR-dialog: A diagnostic dataset for multi-round reasoning in visual dialog. 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)*, pp. 582–595, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1058. URL <https://www.aclweb.org/anthology/N19-1058>.
- Angeliki Lazaridou, Nghia The Pham, and Marco Baroni. Combining language and vision with a multi-modal skip-gram model. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 153–163, Denver, Colorado, May–June 2015. Association for Computational Linguistics. doi: 10.3115/v1/N15-1016. URL <https://www.aclweb.org/anthology/N15-1016>.
- Paul Pu Liang, Zhun Liu, Yao-Hung Hubert Tsai, Qibin Zhao, Ruslan Salakhutdinov, and Louis-Philippe Morency. Learning representations from imperfect time series data via tensor rank regularization.
- Paul Pu Liang, Ziyin Liu, Amir Zadeh, and Louis-Philippe Morency. Multimodal language analysis with recurrent multistage fusion. *arXiv preprint arXiv:1808.03920*, 2018.
- Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Efficient low-rank multimodal fusion with modality-specific factors. *arXiv preprint arXiv:1806.00064*, 2018.
- Louis-Philippe Morency, Ariadna Quattoni, and Trevor Darrell. Latent-dynamic discriminative models for continuous gesture recognition. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*, pp. 1–8. IEEE, 2007.
- Duy-Kien Nguyen and Takayuki Okatani. Multi-task learning of hierarchical vision-language representation. *CoRR*, abs/1812.00500, 2018. URL <http://arxiv.org/abs/1812.00500>.
- Behnaz Nojavanasghari, Deepak Gopinath, Jayanth Koushik, Tadas Baltrušaitis, and Louis-Philippe Morency. Deep multimodal fusion for persuasiveness prediction. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, ICMI 2016, pp. 284–288, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4556-9. doi: 10.1145/2993148.2993176. URL <http://doi.acm.org/10.1145/2993148.2993176>.
- Sunghyun Park, Han Suk Shim, Moitrey Chatterjee, Kenji Sagae, and Louis-Philippe Morency. Computational analysis of persuasiveness in social multimedia: A novel dataset and multimodal prediction approach. In *Proceedings of the 16th International Conference on Multimodal Interaction*, ICMI ’14, pp. 50–57, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2885-2. doi: 10.1145/2663204.2663260. URL <http://doi.acm.org/10.1145/2663204.2663260>.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, and Alexander Ku. Image transformer. *CoRR*, abs/1802.05751, 2018. URL <http://arxiv.org/abs/1802.05751>.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. 2014.
- Ariadna Quattoni, Sybor Wang, Louis-Philippe Morency, Michael Collins, and Trevor Darrell. Hidden conditional random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(10):1848–1852, October 2007. ISSN 0162-8828. doi: 10.1109/TPAMI.2007.1124. URL <http://dx.doi.org/10.1109/TPAMI.2007.1124>.
- Shyam Sundar Rajagopalan, Louis-Philippe Morency, Tadas Baltrušaitis, and Roland Goecke. Extending long short-term memory for multi-view structured learning. In *European Conference on Computer Vision*, 2016.
- Cees G. M. Snoek, Marcel Worring, and Arnold W. M. Smeulders. Early versus late fusion in semantic video analysis. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, MULTIMEDIA ’05, pp. 399–402, New York, NY, USA, 2005. ACM. ISBN 1-59593-044-2. doi: 10.1145/1101149.1101236. URL <http://doi.acm.org/10.1145/1101149.1101236>.

- Yale Song, Louis-Philippe Morency, and Randall Davis. Multi-view latent variable discriminative models for action recognition. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pp. 2120–2127. IEEE, 2012.
- Yale Song, Louis-Philippe Morency, and Randall Davis. Action recognition by hierarchical sequence summarization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3562–3569, 2013.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. *arXiv preprint arXiv:1906.00295*, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Valentin Vielzeuf, Stéphane Pateux, and Frédéric Jurie. Temporal multimodal fusion for video emotion classification in the wild. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction, ICMI '17*, pp. 569–576, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-5543-8. doi: 10.1145/3136755.3143011. URL <http://doi.acm.org/10.1145/3136755.3143011>.
- Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. *arXiv preprint arXiv:1811.09362*, 2018.
- Hao Wu, Jiayuan Mao, Yufeng Zhang, Yuning Jiang, Lei Li, Weiwei Sun, and Wei-Ying Ma. Unified visual-semantic embeddings: Bridging vision and language with structured meaning representations. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Chang Xu, Dacheng Tao, and Chao Xu. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*, 2013.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, abs/1906.08237, 2019. URL <http://arxiv.org/abs/1906.08237>.
- Jiahong Yuan and Mark Liberman. Speaker identification on the scotus corpus. *Journal of the Acoustical Society of America*, 123(5):3878, 2008.
- Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. Mosi: Multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos. *arXiv preprint arXiv:1606.06259*, 2016.
- Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Tensor fusion network for multimodal sentiment analysis. In *Empirical Methods in Natural Language Processing, EMNLP*, 2017.
- Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Memory fusion network for multi-view sequential learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018a.
- Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, and Louis-Philippe Morency. Multi-attention recurrent network for human communication comprehension. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018b.
- AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2236–2246, 2018c.

## A APPENDIX

### A.1 TRAINING REMARKS AND HYPERPARAMETER SPACE SEARCH

The hyperparameters of FMT include the Adam (Kingma & Ba, 2014) learning rate ( $\{0.001, 0.0001\}$ ), structure of summarization network (randomly picked 5 architectures from  $\{1, 2, 3\}$  layers of conv, with kernel shapes of  $\{2, 5, 10, 15, 20\}$ ), number of MTL layers ( $\{4, 6, 8\}$  except for ablation experiments which was  $2 \dots 8$ ), number of FMT units ( $\{4, 6\}$ , except for ablation experiment which was  $1 \dots 6$ ),  $e_{M \in \{L, V, A\}}$  ( $\{20, 40\}$ ), dropout (0, 0.1). The same parameters (when applicable) are used for training MulT for POM dataset (e.g. num encoder layers same as number of MTL). Furthermore, for MulT specific hyperparameters, we use similar values as Table 5 in the original paper. All models are trained for a maximum of 200 epochs. The hyperparameter validation is similar to Zadeh et al. (2018b).

### A.2 NUMBER OF MTL

We study the effect of number of MTL on FMT performance. Table 6 shows the results of this experiment. The best performance is achieved using 8 MTL layers (which was also the maximum layers we tried in our hyperparameter search).

### A.3 NUMBER OF ATTENTION HEADS FOR ORIGINAL TRANSFORMER MODEL

In this section, we discuss the effect of increasing the number of heads on the original transformer model (OTF, Vaswani et al. (2017)). Please note that we implement the OTF to allow for all attention heads to have full input receptive field (from  $1 \dots \mathcal{T}$ ), similar to FMT. We increase the attention heads from 1 to 35 (after 35 does not fit on a Tesla-V100 GPU with batchsize of 20). Table 7 shows the results of increasing number of attention heads for both models. We observe that achieving superior performance is not a matter of increasing the attention heads. Even using 1 FMS unit, which leads to 7 total attention, FMT achieves higher performance than counterpart OTF.

### A.4 TRAINING REMARKS FOR MORE THAN 3 MODALITIES

In many scenarios in nature, as well as what is currently pursued in machine learning, the number of modalities goes as high as 3 (mostly language, vision and acoustic, as studied in this paper). This leads to 7 attentions within each FMS, well manageable for successful training of FMT as demonstrated in this paper. However, as the number of modalities increases, the underlying multimodal phenomena becomes more challenging to model. This causes complexities for any competitive multimodal model, regardless of their internal design. While studying these cases are beyond the scope of this paper, due to rare nature of having more than 3 main modalities modalities, for FMT, the complexity can be managed due to the factorization in FMS. We propose two approaches: 1) for high number of modalities, the involved factors can be reduced based on domain knowledge, the nature of the problem, and the assumed dependencies between modalities (e.g. removing factors between modalities that are deemed weakly related). Alternatively, without making assumptions about

<i>Model \ Metric</i>	<i>MAE</i>	<i>Corr</i>
FMT [2]	0.881	0.720
FMT [3]	0.876	0.727
FMT [4]	0.871	0.723
FMT [5]	0.876	0.724
FMT [6]	0.852	0.730
FMT [7]	0.859	0.732
FMT [8]	<b>0.837</b>	<b>0.744</b>

Table 6: Multimodal sentiment analysis experiments with different number of MTL layers. The number in the square bracket indicates the number of MTL layers.

<i>Model \ Metric</i>	<i>BA</i>	<i>F1</i>	<i>MAE</i>	<i>Corr</i>
OTF [1]	74.6/76.5	74.5/76.5	0.983	0.651
OTF [2]	76.8/78.8	76.6/78.9	0.975	0.655
OTF [3]	74.2/75.8	74.0/75.9	0.998	0.647
OTF [4]	76.5/77.9	76.6/78.2	1.022	0.632
OTF [5]	75.1/76.7	75.1/76.6	1.026	0.626
OTF [6]	71.6/72.3	71.3/72.6	1.094	0.677
OTF [7]	77.4/79.1	77.4/79.0	0.988	0.646
OTF [14]	77.0/78.5	76.9/78.5	0.972	0.683
OTF [21]	75.9/77.7	75.8/77.9	0.930	0.682
OTF [35]	67.6/68.9	67.2/73.2	1.174	0.502

Table 7: Results of experiments with different number of heads for OTF. The number in the square bracket indicates the number of heads.

inter-modality dependencies, a greedy approach may be taken for adding factors; an approach similar to stepwise regression (Kleinbaum et al., 1988), iteratively adding the next most important factor. Using these two methods, the model can cope with higher number of modalities with a controllable compromise between performance and overparameterization.