

Learning Relationships between Text, Audio, and Video via Deep Canonical Correlation for Multimodal Language Analysis

Zhongkai Sun,¹ Prathusha K Sarma,^{2*} William A Sethares,¹ Yingyu Liang¹

¹University of Wisconsin-Madison, ²Curai
{zsun227, sethares}@wisc.edu, prathyushaks.21@gmail.com, yliang@cs.wisc.edu

Abstract

Multimodal language analysis often considers relationships between features based on text and those based on acoustical and visual properties. Text features typically outperform non-text features in sentiment analysis or emotion recognition tasks in part because the text features are derived from advanced language models or word embeddings trained on massive data sources while audio and video features are human-engineered and comparatively underdeveloped. Given that the text, audio, and video are describing the same utterance in different ways, we hypothesize that the multimodal sentiment analysis and emotion recognition can be improved by learning (hidden) correlations between features extracted from the outer product of text and audio (we call this *text-based audio*) and analogous *text-based video*. This paper proposes a novel model, the Interaction Canonical Correlation Network (ICCN), to learn such multimodal embeddings. ICCN learns correlations between all three modes via deep canonical correlation analysis (DCCA) and the proposed embeddings are then tested on several benchmark datasets and against other state-of-the-art multimodal embedding algorithms. Empirical results and ablation studies confirm the effectiveness of ICCN in capturing useful information from all three views.

1 Introduction

Human language communication occurs in several modalities: via words that are spoken, by tone of voice, and by facial and bodily expressions. Understanding the content of a message thus requires understanding all three modes. With the explosive growth in availability of data, several machine learning algorithms have been successfully applied towards multimodal tasks such as sentiment analysis (Morency, Mihalcea, and Doshi 2011; Soleymani et al. 2017), emotion recognition (Haq and Jackson 2011), image-text retrieval (Wang, Li, and Lazebnik 2016), and aiding medical diagnose (Liu et al. 2019; Lee et al. 2014) etc. Among multimodal language sentiment or emotion experiments involving unimodal features (Zadeh et al. 2016; 2018c; Tsai et al. 2018; 2019), it is commonly observed that text based features perform better than visual or auditory

modes. This is plausible for at least three reasons: 1) Text itself contains considerable sentiment-related information. 2) Visual or acoustic information may sometimes confuse the sentiment or emotion analysis task. For instance: “angry” and “excited” may have similar acoustic performances (high volume and high pitch) even though they belong to opposite sentiments. Similarly, “sad” and “disgusted” may have different visual features though they both belong to the negative sentiment. 3) Algorithms for text analysis have a richer history and are well studied.

Based on this observation, learning the hidden relationship between verbal information and non-verbal information is a key point in multi-modal language analysis. This can be approached by looking at different ways of combining multimodal features.

The simplest way to combine text (T), audio (A) and video (V) for feature extraction and classification is to concatenate the A, V, and T vectors. An alternative is to use the outer product (Liu et al. 2018; Zadeh et al. 2017) which can represent the interaction between pairs of features, resulting in 2D or 3D arrays that can be processed using advanced methods such as convolutional neural networks (CNNs) (Lawrence et al. 1997). Other approaches (Zadeh et al. 2018a; Liang et al. 2018; Zadeh et al. 2018c; Wang et al. 2018) study multi-modal interactions and intra-actions by using either graph or temporal memory networks with a sequential neural network LSTM (Gers, Schmidhuber, and Cummins 1999). While all these have contributed towards learning multi-modal features, they typically ignore the hidden correlation between text-based audio and text-based video. Individual modalities are either combined via neural networks or passed directly to the final classifier stage. However, it is obvious that attaching both audio and video features to the same textual information may enable non-text information to be better understood, and in turn the non-text information may impart greater meaning to the text. Thus, it is reasonable to study the deeper correlations between text-based audio and text-based video features.

This paper proposes a novel model which uses the outer-product of feature pairs along with Deep Canonical Correlation Analysis (DCCA) (Andrew et al. 2013) to study useful multi-modal embedding features. The effectiveness of us-

*Work done while at UW-Madison

ing an outer-product to extract cross-modal information has been explored in (Zadeh et al. 2017; Liu et al. 2018). Thus, features from each mode are first extracted independently at the sentence (or utterance) level and two outer-product matrices ($T \otimes V$ and $T \otimes A$) are built for representing the interactions between text-video and between text-audio. Each outer-product matrix is connected to a convolutional neural network (CNN) for feature extraction. Outputs of these two CNNs can be considered as feature vectors for text-based audio and text-based video and should be correlated.

In order to better correlate the above text-based audio and text-based video, we use Canonical Correlation Analysis (CCA) (Hotelling 1936), which is a well-known method for finding a linear subspace where two inputs are maximally correlated. Unlike cosine similarity or Euclidean distance, CCA is able to learn the direction of maximum correlation over all possible linear transformations and is not limited by the original coordinate systems. However, one limitation of CCA is that it can only learn linear transformations. An extension to CCA named Deep CCA (DCCA) (Andrew et al. 2013) uses a deep neural network to allow non-linear relationships in the CCA transformation. Recently several authors (Rotman, Vulić, and Reichart 2018; Hazarika et al. 2018) have shown the advantage of using CCA-based methods for studying correlations between different inputs. Inspired by these, we use DCCA to correlate text-based audio and text-based video. Text-based audio and text-based video features derived from the two CNNs are input into a CCA layer which consists of two projections and a CCA Loss calculator. The two CNNs and the CCA layer then form a DCCA, the weights of the two CNNs and the projections are updated by minimizing the CCA Loss. In this way, the two CNNs are able to extract useful features from the outer-product matrices constrained by the CCA loss. After optimizing the whole network, outputs of the two CNNs are concatenated with the original text sentence embedding as the final multi-modal embedding, which can be used for the classification.

We evaluate our approach on three benchmark multi-modal sentiment analysis and emotion recognition datasets: CMU-MOSI (Zadeh et al. 2016), CMU-MOSEI (Zadeh et al. 2018c), and IEMOCAP (Busso et al. 2008). Additional experiments are presented to illustrate the performance of the ICCN algorithm. The rest of the paper is organized as follows: Section 2 presents related work, Section 3 introduces our proposed model and Section 4 describes our experimental setup. Section 5 presents a discussion on the empirical observations, Section 6 concludes this work.

2 Related Work

The central themes of this paper are related to learning (i) multi-modal fusion embeddings and (ii) cross-modal relationships via canonical correlation analysis (CCA).

Multi-modal fusion embedding: Early work (Poria et al. 2016) concatenates the audio, video and text embeddings to learn a larger multi-modal embedding. But this may lead to a potential loss of information between different modalities. Recent studies on learning multi-modal fusion embeddings train specific neural network architectures to combine

all three modalities. In their work (Chen et al. 2017) propose improvements to multi-modal embeddings using reinforcement learning to align the multi-modal embedding at the word level by removing noises. A multi-modal tensor fusion network is built in (Zadeh et al. 2017) by calculating the outer-product of text, audio and video features to represent comprehensive features. However this method is limited by the need of a large computational resources to perform calculations of the outer dot product. In their work (Liu et al. 2018) developed an efficient low rank method for building tensor networks which reduce computational complexity and are able to achieve competitive results. A Memory Fusion Network (MFN) is proposed by (Zadeh et al. 2018a) which memorizes temporal and long-term interactions and intra-actions between cross-modals, this memory can be stored and updated in a LSTM. (Liang et al. 2018) learned multistage fusion at each LSTM step so that the multi-modal fusion can be decomposed into several sub-problems and then solved in a specialized and effective way. A multimodal transformer is proposed by (Tsai et al. 2019) that uses attention based cross-modal transformers to learn interactions between modalities.

Cross-modal relationship learning via CCA: Canonical Correlation Analysis (CCA) (Hotelling 1936) learns the maximum correlation between two variables by mapping them into a new subspace. Deep CCA (DCCA) (Andrew et al. 2013) improves the performance of CCA by using feed-forward neural networks in place of the linear transformation in CCA.

A survey of recent literature sees applications of CCA-based methods in analyzing the potential relationship between different variables. For example, a CCA based model to combine domain knowledge and universal word embeddings is proposed by (Sarma, Liang, and Sethares 2018). Models proposed by (Rotman, Vulić, and Reichart 2018) use Deep Partial Canonical Correlation Analysis (DPCCA), a variant of DCCA, for studying the relationship between two languages based on the same image they are describing. Work by (Sun et al. 2019) investigates the application of DCCA to simple concatenations of multimodal-features, while (Hazarika et al. 2018) applied CCA methods to learn joint-representation for detecting sarcasm. Both approaches show the effectiveness of CCA methods towards learning potential correlation between two input variables.

3 Methodology

This section first introduces CCA and DCCA. Next, the interaction canonical correlation network (ICCN), which extracts the interaction features of a CNN in a DCCA-based network, is introduced. Finally, the whole pipe-line of using this method in a multimodal language analysis task is described.

CCA and DCCA

Given two sets of vectors $X \in \mathbb{R}^{n_1 \times m}$ and $Y \in \mathbb{R}^{n_2 \times m}$, where m denotes the number of vectors, CCA learns two linear transformations $A \in \mathbb{R}^{n_1 \times r}$ and $B \in \mathbb{R}^{n_2 \times r}$ such that the correlation between $A^T X$ and $B^T Y$ is maximized. De-

note the covariances of X and Y as S_{11}, S_{22} , and the cross-covariance of X, Y as S_{12} . The CCA objective is

$$\begin{aligned} A^*, B^* &= \arg \max_{A, B} \text{corr}(A^T X, B^T Y) \\ &= \arg \max_{A, B} \frac{A^T S_{12} B}{\sqrt{A^T S_{11} A} \sqrt{B^T S_{22} B}}. \end{aligned} \quad (1)$$

The solution of the above equation is fixed and can be solved in multiple ways (Hotelling 1936; Martin and Maes 1979). One method suggested by (Martin and Maes 1979) lets U, S, V^T be the Singular Value Decomposition (SVD) of the matrix $Z = S_{11}^{-\frac{1}{2}} S_{12} S_{22}^{-\frac{1}{2}}$. Then A^*, B^* and the total maximum canonical correlation are

$$\begin{aligned} A^* &= S_{11}^{-\frac{1}{2}} U \\ B^* &= S_{22}^{-\frac{1}{2}} V \\ \text{corr}(A^{*T} X, B^{*T} Y) &= \text{trace}(Z^T Z)^{\frac{1}{2}}. \end{aligned} \quad (2)$$

One limitation of CCA is that it only considers linear transformations. DCCA (Andrew et al. 2013) learns non-linear transformations using a pair of neural networks. Let f, g denote two independent neural networks, the objective of DCCA is to optimize parameters θ_f, θ_g of f, g so that the canonical correlation between the output of f and g , denoted as $F_X = f(X; \theta_1)$ and $F_Y = g(Y; \theta_2)$, can be maximized by finding two linear transformations C^*, D^* . The objective of DCCA is

$$\begin{aligned} \theta_f^*, \theta_g^* &= \arg \max_{\theta_f, \theta_g} \text{CCA}(F_X, F_Y) \\ &= \arg \max_{\theta_f, \theta_g} \text{corr}(C^{*T} F_X, D^{*T} F_Y). \end{aligned} \quad (3)$$

In order to update the parameters of f, g , a loss for measuring the canonical correlation must be calculated and back-propagated. Let R_{11}, R_{22} be covariances of F_X, F_Y , the cross-covariance of F_X, F_Y as R_{12} . Let $E = R_{11}^{-\frac{1}{2}} R_{12} R_{22}^{-\frac{1}{2}}$. According to (2), the canonical correlation loss for updating F_x, F_Y can be defined by

$$\text{CCA Loss} = -\text{trace}(E^T E)^{\frac{1}{2}}. \quad (4)$$

Networks $f(X; \theta_f), g(Y; \theta_g)$'s parameters can be updated by minimizing the CCA Loss (4) (i.e. maximize the total canonical correlation).

Text Based Audio Video Interaction Canonical Correlation

Previous work of (Zadeh et al. 2017; Liu et al. 2018) on multi-modal feature fusion has shown that the outer-product is able to learn interactions between different features effectively. Thus, we use the outer-product to represent text-video and text-audio features. Given that outer-product and DCCA are applied at the utterance (sentence)-level, we extract utterance level features for each uni-modal independently in order to test the effectiveness of ICCN more directly. Let $H_t \in \mathbb{R}^{d_t}$ be the utterance-level text feature embedding, and $H_v \in \mathbb{R}^{d_v \times l_v}, H_a \in \mathbb{R}^{d_a \times l_a}$ be the video

and audio input sequences. A 1D temporary convolutional layer is used to extract local structure of the audio and video sequences, and the outputs of the 1D-CNN are denoted as $H_{a1} \in \mathbb{R}^{d_{a1} \times l_a}, H_{v1} \in \mathbb{R}^{d_{v1} \times l_v}$. Next, two LSTMs process the audio and video sequences. The final hidden state of each LSTM is used as the utterance-level audio or video feature, denoted as $H_{a2} \in \mathbb{R}^{d_{a2}}, H_{v2} \in \mathbb{R}^{d_{v2}}$. Once each utterance-level feature has been obtained, the text-based audio feature matrix and text-based video feature matrix can then be learned using the outer-product on H_t, H_{v2}, H_{a2} :

$$\begin{aligned} H_{ta} &= H_t \otimes H_{a2}, H_{ta} \in \mathbb{R}^{d_t \times d_{a2}} \\ H_{tv} &= H_t \otimes H_{v2}, H_{tv} \in \mathbb{R}^{d_t \times d_{v2}}. \end{aligned} \quad (5)$$

In order to extract useful features from the outer-product matrices H_{ta}, H_{tv} , a Convolutional Neural Network is used as the basic feature extractor. H_{ta} and H_{tv} are connected by multiple 2D-CNN layers with max pooling. Outputs of the two 2D-CNNs are reshaped to 1D vector and then be used as inputs to the CCA Loss calculation. 1D-CNN, LSTM, and 2D-CNN's weights are again updated using the back-propagation of the CCA Loss (4). Thus the two 2D-CNNs learn to extract features from H_{tv} and H_{ta} so that their canonical correlation is maximized. Algorithm 1 provides the pseudo-code for the whole Interaction Canonical Correlation Network (ICCN).

Algorithm 1 Interaction Canonical Correlation Network

Input: Data $H_t \in \mathbb{R}^{d_t \times N}, H_a \in \mathbb{R}^{d_a \times l_a \times N}, H_v \in \mathbb{R}^{d_v \times l_v \times N}, epoch, \eta$
Initialize W_a of (CNN1D_a, LSTM_a, and CNN2D_{ta})
Initialize W_v of (CNN1D_v, LSTM_v, and CNN2D_{tv})
while $epoch > 0$ **do**
 $H_{a1} = \text{CNN1D}_a(H_a)$
 $H_{v1} = \text{CNN1D}_v(H_v)$
 $H_{a2} = \text{LSTM}_a(H_{a1})$
 $H_{v2} = \text{LSTM}_v(H_{v1})$
 $H_{tv} = H_t \otimes H_{v2}, H_{ta} = H_t \otimes H_{a2}$
 $K_{tv} = \text{CNN2D}_{tv}(H_{tv})$
 $K_{ta} = \text{CNN2D}_{ta}(H_{ta})$
Compute gradients $\nabla W_v, \nabla W_a$ of:
 $\min_{W_v, W_a} \text{CCALoss}(K_{tv}, K_{ta})$
Update:
 $W_v \leftarrow W_v - \eta \nabla W_v$
 $W_a \leftarrow W_a - \eta \nabla W_a$
 $epoch \leftarrow epoch - 1$
end while
Output: K_{tv}, K_{ta}

Pipe-line for Downstream Tasks

The ICCN method acts as a feature extractor. In order to test its performance, an additional downstream classifier is also required. Uni-modal features can be extracted using a variety of simple extraction schemes as well as by learning features using more complex neural network based models such as a sequential LSTM. Once uni-modal features for

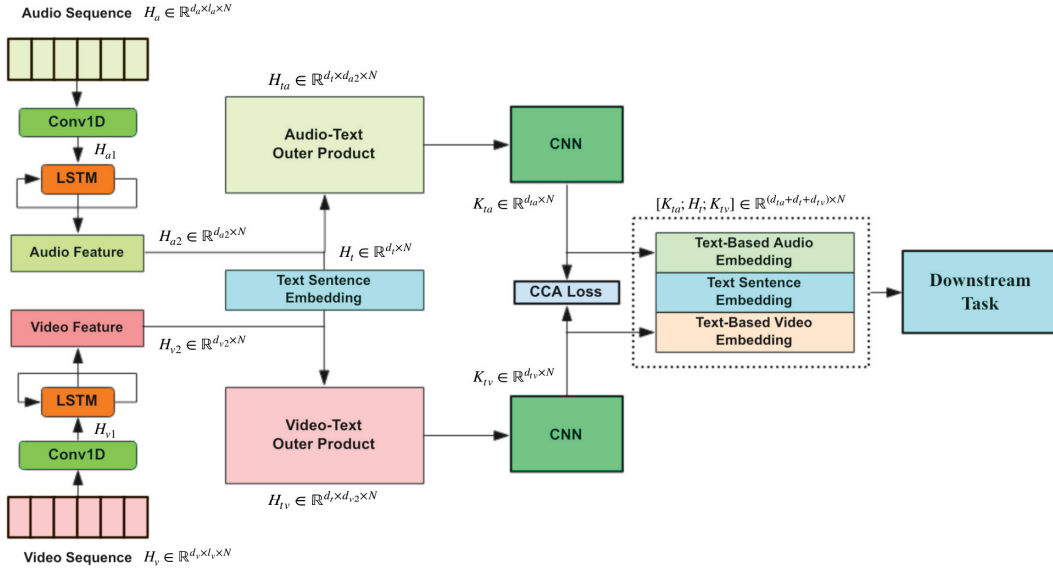


Figure 1: ICCN method for aligning text-based audio features and text-based video features. Sentence level uni-modal features are extracted independently. Outer-product matrices of text-audio and text-video are used as input to the Deep CCA network. After learning the CNN’s weights using the CCA Loss, outputs of the two CNNs are concatenated with the original text to form the multi-modal embedding. This can be used as input to independent downstream tasks.

text, video, and audio have been obtained, the ICCN can be used to learn text-based audio feature K_{ta} and text-based video feature K_{tv} . The final multi-modal embedding can be formed as the concatenation of the text-based audio, the original text, and the text-based video features, which are denoted as $[K_{ta}; H_t; K_{tv}]$. This $[K_{ta}; H_t; K_{tv}]$ can then be used as an input to downstream classifiers such as logistic regression or multilayer perceptron. Figure 1 shows the pipeline using the ICCN for downstream tasks in our work.

4 Experiment Settings

This section describes the experimental datasets and baseline algorithms against which ICCN is compared.

Datasets

The proposed algorithm is tested using three public benchmark multi-modal sentiment analysis and emotion recognition datasets: CMU-MOSI (Zadeh et al. 2016), CMU-MOSEI (Zadeh et al. 2018c), and IEMOCAP (Busso et al. 2008). Both CMU-MOSI and CMU-MOSEI’s raw features, and most of the corresponding extracted features can be acquired from the CMU-MultimodalSDK (Zadeh et al. 2018b).

- **CMU-MOSI:** This dataset is a multi-modal dataset built on 93 Youtube movie reviews. Videos are segmented to 2198 utterance clips, and each utterance example is annotated on a scale of $[-3, 3]$ to reflect sentiment intensity. -3 means an extremely negative and 3 means an extremely positive sentiment. This data set is divided into three parts, training (1283 samples), validation (229 samples) and test (686 samples).

- **CMU-MOSEI:** This dataset is similar to the CMU-MOSI, but is larger in size. It consists of 22856 annotated utterances extracted from Youtube videos. Each utterance can be treated as an individual multi-modal example. Train, validation, and test sets contain 16326, 1871, and 4659 samples respectively.
- **IEMOCAP:** This dataset contains 302 videos in which speakers performed 9 different emotions (angry, excited, fear, sad, surprised, frustrated, happy, disappointed and neutral). Those videos are divided into short segments with emotion annotations. Due to the imbalance of some emotion labels, we follow experiments in previous papers (Wang et al. 2018; Liu et al. 2018) where only four emotions (angry, sad, happy, and neutral) are used to test the performance of the algorithm. Train, validation, and test partitions contain 2717, 789, and 938 data samples respectively.

Multi-modal Features

The following uni-modal features are used prior to combinations,

- **Text Features:** For MOSI and MOSEI, we use a pre-trained transformer model BERT (Devlin et al. 2018) to extract utterance level text features, (many other approaches use Glove word-level embeddings followed by a LSTM). The motivation behinds using BERT is 1) BERT is the state-of-the-art in sentence encoding algorithms and has demonstrated tremendous success in several downstream text applications such as sentiment analysis, question-answering, semantic similarity tasks etc, 2) using BERT simplifies the training pipe-line, with a large

focus now towards improving the performance of ICCN on a particular downstream task. We input the raw text to the pre-trained uncased BERT-Base model (without fine-tuning). Sentence encodings output from BERT are used as the text features. Each individual text feature is of size 768. For IEMOCAP, we used InferSent (Conneau et al. 2017) to encode utterance level text. Since data is provided in the form of word indices for GLOVE embeddings rather than raw text, we use InferSent; a BiLSTM layer followed by a max-pooling layer to learn sentence embeddings.

- **Audio Features:** The audio feature is extracted by using COVAREP (Degottex et al. 2014), which is a public software used for extracting acoustic features such as pitch, volume, and frequency. The CMU-MultimodalSDK provides COVAREP feature sequence for every multi-modal example, the dimension of each frame’s audio feature is 74.
- **Video Features:** Facet (iMotions 2017) has been used for extracting facial expression features such as action units and face pose. Similarly, every multimodal example’s video feature sequence is also obtained from the CMU-MultimodalSDK. The size of each frame’s video feature is 35.

Baseline Methods

We consider a variety of baseline methods for multi-modal embedding comparison. In order to focus on the multi-modal embedding itself, we input each multi-modal embedding to the same downstream task classifier (or regressor). Experimental comparisons are reported in two parts 1) we report the effectiveness of DCCA over the simpler CCA based methods when used as inputs to the ICCN and 2) we compare ICCN against newer utterance level embeddings algorithms that learn features for a down stream task in an end-to-end fashion. The following baselines are used in our comparisons,

- **Uni-modal and Concatenation:** This is the simplest baseline in which uni-modal features are concatenated to obtain a multi-modal embedding.
- **Linear CCA:** CCA (Hotelling 1936) considers linear transformations for different inputs. We use the CCA to learn a new common space for audio and video modes, and combine the learned audio and video features with the original text embedding. This is because, (Sun et al. 2019) showed that using a CCA-based method to correlate audio and video is more effective than correlating audio-text or video-text.
- **Kernel-CCA:** Kernel-CCA (Akaho 2006) introduces a nonlinearity via kernel maps. Kernel-CCA can be used exactly like CCA.
- **GCCA:** Generalized CCA (Tenenhaus and Tenenhaus 2011) learns a common subspace across multiple views. We use GCCA in two ways: 1) use the GCCA output embedding directly and 2) combine the GCCA output embedding with the original text embedding.

- **DCCA:** A Deep CCA based algorithm proposed by (Sun et al. 2019). Audio and video features are simply concatenated and then be correlated with text features using DCCA. Outputs of the DCCA are again concatenated with raw text, audio, and video features to formulate the multimodal embedding.

In the proposed ICCN algorithm, text features are encoded by a pre-trained BERT transformer or by InferSent. This is unlike most of the state-of-the-art algorithms that obtain sentence level encodings by passing word embeddings through variants of RNNs. However, since the idea is to compare modal features, we also choose the following three state-of-the-art utterance-level (i.e. sentence-level) fusion models (whose core algorithm is agnostic to the text encoding architecture) as additional baselines. To make the comparison fair, these methods use the same features as ours.

- **TFN:** Tensor Fusion Network (TFN) (Zadeh et al. 2017) combines individual modal’s embeddings via calculating three different outer-product sub-tensors: unimodal, bimodal, and trimodal. All tensors will then be flattened and used as a multi-modal embedding vector.
- **LMF:** Low-Rank Multimodal Fusion (LMF) (Liu et al. 2018) learns the multimodal embedding based the similar tensor processing of TFN, but with an additional low-rank factor for reducing computation memory.
- **MFM:** Multimodal Factorization Model (MFM) (Tsai et al. 2018) is consists of a discriminative model for prediction and a generative model for reconstructing input data. A comprehensive multimodal embedding is learned via optimizing the generative-discriminative objective simultaneously.

Ablation studies of ICCN

In order to analyze the usefulness of different components of the ICCN, we consider the following two questions:

Question 1: Is using canonical correlation better than using other methods like Cosine-Similarity?

Question 2: Is learning the interactions between text and video (or audio) useful?

We design several experiments to address these two questions, First, we replace the CCA Loss with Cosine-Similarity Loss while leaving other parts of the ICCN unchanged. Second, instead of using the outer-product of audio (or video) and text as input to the CCA Loss, we use audio and video directly as the input to DCCA network. We compare different ICCN variants’ performance to prove the usefulness of each component of the network.

Evaluation Methods

To evaluate ICCN against previous baselines, the following performance metrics as described in (Liu et al. 2018; Zadeh et al. 2017; Tsai et al. 2018) are reported,

- On the CMU-MOSI and CMU-MOSEI we report four performance metrics, i) binary accuracy, ii) F1-score iii) mean absolute error and iv) 7-class sentiment level / Correlation with human labeling.

Data View	CMU-MOSI					CMU-MOSEI				
	Acc-2	F-score	MAE	Acc-7	Corr	Acc-2	F-score	MAE	Acc-7	Corr
Audio	45.15	45.83	1.430	16.21	0.248	58.75	59.23	0.785	38.59	0.298
Video	48.10	49.06	1.456	15.51	0.339	59.25	59.90	0.770	36.09	0.288
Text	80.80	80.17	0.897	35.92	0.688	82.83	83.02	0.582	48.76	0.681
Text+Video	81.00	80.91	0.920	35.11	0.676	82.86	83.01	0.581	47.92	0.674
Text+Audio	80.59	80.56	0.909	35.08	0.672	82.80	82.96	0.582	49.02	0.689
Audio+Video+Text	80.94	81.00	0.895	36.41	0.689	82.72	82.87	0.583	50.11	0.692
CCA	79.45	79.35	0.893	34.15	0.690	82.94	83.06	0.573	50.23	0.690
KCCA	79.82	79.91	0.889	34.76	0.689	83.05	83.14	0.574	50.09	0.692
GCCA	62.50	62.15	1.403	17.29	0.533	75.12	75.46	0.653	45.33	0.602
GCCA+Text	77.80	77.87	1.107	25.94	0.658	82.75	82.90	0.613	46.06	0.644
DCCA	80.60*	80.57*	0.874	35.51	0.703	83.62*	83.75*	0.579	50.12	0.707
TFN	80.82	80.77	0.901	34.94	0.698	82.57	82.09	0.593	50.21	0.700
LMF	82.53	82.47	0.917	33.23	0.695	82.03	82.18	0.623	48.02	0.677
MFN	81.72	81.64	0.877	35.42	0.706	84.40	84.36	0.568	51.37	0.717
ICCN	83.07	83.02	0.862	39.01	0.714	84.18	84.15	0.565	51.58	0.713

Table 1: Results for experiments on CMU-MOSI and CMU-MOSEI. Best numbers are in bold. For accuracy, F-score, and Correlation, higher is better. For mean absolute error, lower is better. Results marked with * are reported in original papers. For TFN, LMF, and MFN, we re-did experiments with using our features for a fair comparison.

- On the IEMOCAP we used i) binary accuracy and ii) F1-score for evaluation.

Evaluations Details on CMU-MOSI, CMU-MOSEI:

The original MOSI and MOSEI datasets are labeled in the range [-3,3]. The author of the datasets suggests a criterion for building binary labels: examples with label in [-3, 0) are considered to have negative sentiment while examples with label in (0, 3] are considered to have positive sentiment. 7-class sentiment level is also calculated based on the label distribution in [-3,3]. The correlation of predicted results with human labeling is also used as a criteria.

Hyperparameter Tuning

A basic Grid-Search is used to tune hyperparameters, and the best hyperparameter settings for the ICCN are chosen according to its performance on the validation dataset. For ICCN, hyperparameters and tuning ranges are: learning rate ($1e - 5$ – $1e - 3$), mini-batch size (128–512), the number of epoch (10–100), hidden dimensions of MLP (64–512), and output dimension of the CCA projection (30–100). ReLU is used as the activation function, RMSProp is used as the optimizer.

Whenever the training of the ICCN with a specific hyperparameter setting has finished, features learned from the ICCN are used as input to the same downstream task models (a simple MLP is used in this work). Test results are reported by using the best hyperparameter setting learned above.

5 Discussion of Empirical Results

This section presents and discusses results on the CMU-MOSI, CMU-MOSEI, and IEMOCAP datasets.

Performance on Benchmark Datasets

Tables 1 and 2 present results of experiments on the CMU-MOSI, CMU-MOSEI, and IEMOCAP datasets.

- First, when compared with results of using uni-modal and simple concatenations, ICCN outperforms all of them in all of the criteria. Note that the performance of the text feature is always better than the performance of the audio

and video, and that a simple concatenation of text, video, and audio does not work well. This shows the advance of highly-developed pre-trained text features capable of improving the overall multimodal task performance. However, it also shows the challenge of how to effectively combine such a highly developed text feature with audio and video features.

- Second, ICCN also outperforms other CCA-based methods. The results of using other CCA-based methods show that they cannot improve the multimodal embedding’s performance. We argue this occurs because 1) CCA / KCCA / GCCA do not exploit the power of neural network architectures so that their learning capacities are limited. 2) Using DCCA without learning the interactions between text-based audio and text-based video may sacrifice useful information.
- Third, the ICCN still achieves better or similar results when compared with other neural network based state-of-the-art methods (TFN, LMF, and MFN). These results demonstrate the ICCN’s competitive performance.

Results of Ablation Studies

Table 3 shows results of using variants of ICCN on CMU-MOSI and CMU-MOSEI datasets.

First, using the CCA Loss performs better than using Cosine-Similarity Loss with or without the outer-product. This is reasonable as the DCCA is able to learn the hidden relationships (with the help of non-linear transformations) but cosine-similarity is restrained by the original coordinates. To further verify this, we also record changes of canonical correlation and cosine similarity between text-based audio and text-based video (i.e., between the two outputs of the CNNs in the ICCN) by using CCA Loss or Cosine Similarity Loss for the ICCN with using the CMU-MOSI dataset. Curves in Figures 2 and 3 summarize the results of the experiments. Results show that maximizing the canonical correlation by using the CCA Loss does not necessarily increase the cosine similarity, and vice versa. This demonstrates that the canonical correlation is a genuinely

Data View Emotions	IEMOCAP							
	Happy		Angry		Sad		Neutral	
	Acc-2	F-score	Acc-2	F-score	Acc-2	F-score	Acc-2	F-score
Audio	84.03	81.09	85.49	84.03	82.75	80.26	63.08	59.24
Video	83.14	80.36	85.91	83.27	81.19	80.35	62.30	58.19
Text	84.80	81.17	85.12	85.21	84.18	83.63	66.02	63.52
Text+Video	85.32	82.01	85.22	85.10	83.33	82.96	65.82	65.81
Text+Audio	85.10	83.47	86.09	84.99	83.90	83.91	66.96	65.02
Audio+Video+Text	86.00	83.37	86.37	85.88	84.02	83.71	66.87	65.93
CCA	85.91	83.32	86.17	84.39	84.19	83.71	67.22	64.84
KCCA	86.54	84.08	86.72	86.32	85.03	84.91	68.29	65.93
GCCA	81.15	80.33	82.47	78.06	83.22	81.75	65.34	59.99
GCCA+Text	87.02	83.44	88.01	88.00	84.79	83.26	68.26	67.61
DCCA	86.99	84.32	87.94	87.85	86.03	84.36	68.87	65.93
TFN	86.66	84.03	87.11	87.03	85.64	85.75	68.90	68.03
LMF	86.14	83.92	86.24	86.41	84.33	84.40	69.62	68.75
MFM	86.67	84.66	86.99	86.72	85.67	85.66	70.26	69.98
ICCN	87.41	84.72	88.62	88.02	86.26	85.93	69.73	68.47

Table 2: Results for experiments on IEMOCAP. Best numbers are in bold. For accuracy and F-score, higher is better. For TFN, LMF, and MFM, we re-did experiments with using our features for a fair comparison.

Data View	CMU-MOSI					CMU-MOSEI				
	Acc-2	F-score	MAE	Acc-7	Corr	Acc-2	F-score	MAE	Acc-7	Corr
ICCN	83.07	83.02	0.862	39.01	0.714	84.18	84.15	0.565	51.58	0.713
ICCN ₁ (no text)	82.13	82.05	0.874	35.51	0.703	83.01	83.10	0.575	50.12	0.707
ICCN ₂ (cos)	82.32	82.27	0.876	36.01	0.702	82.98	82.90	0.575	50.63	0.700
ICCN ₃ (no text + cos)	81.49	81.58	0.889	35.77	0.692	82.59	82.73	0.578	50.21	0.696

Table 3: Ablation studies of ICCN on CMU-MOSI and CMU-MOSEI. ICCN₁₋₃ denote different variants, “no text” means applying DCCA to audio and video directly instead of applying to the outer-product with text; “cos” means replacing CCA Loss with Cosine-Similarity Loss; “no text + cca” means removing outer-product with text and use Cosine-Similarity Loss.

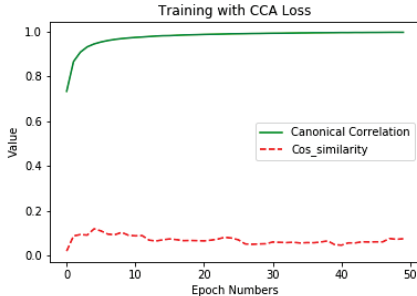


Figure 2: Changes of mean canonical correlation and mean cosine similarity between text-based audio and text-based video when training with CCA Loss: The network learns to maximize the canonical correlation but the cosine similarity isn’t affected.

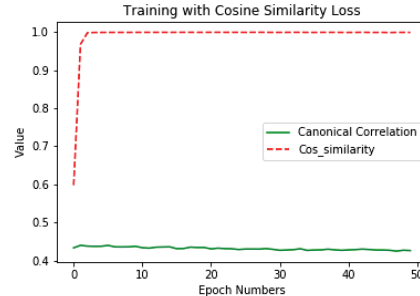


Figure 3: Changes of mean canonical correlation and cosine similarity between text-based audio and text-based video when training with Cosine Similarity Loss: convergence of cosine similarity doesn’t affect canonical correlation.

different objective function than cosine similarity, and explains the different behaviors in the downstream applications. CCA is capable of learning hidden relationships between inputs that the cosine similarity does not see.

Second, learning the interactions between non-text and text performs better than using audio and video directly. This also makes sense because audio and video are more correlated when they are based on the same text, thus learning text-based audio and text-based video performs better. In

summary, Table 3 shows the usefulness of using a text-based outer-product together with DCCA.

6 Conclusion and Future Work

This paper has proposed the ICCN method, which uses canonical correlation to analyze hidden relationships between text, audio, and video. Testing on a multi-modal sentiment analysis and emotion recognition task shows that the multi-modal features learned from the ICCN model can achieve state-of-the-art performance, and shows the effec-

tiveness of the model. Ablation studies confirm the usefulness of different part of the network.

There is, of course, considerable room for improvement. Possible directions include learning dynamic intra-actions in each model together with inter-actions between different modes; learning the trade-off between maximum canonical correlation and best downstream task performance; and developing an interpretable end-to-end multi-modal canonical correlation model. In the future, we hope to move forward in the development of multi-modal machine learning.

References

- Akaho, S. 2006. A kernel method for canonical correlation analysis. *arXiv preprint cs/0609071*.
- Andrew, G.; Arora, R.; Bilmes, J.; and Livescu, K. 2013. Deep canonical correlation analysis. In *ICML*, 1247–1255.
- Busso, C.; Bulut, M.; Lee, C.-C.; Kazemzadeh, A.; Mower, E.; Kim, S.; Chang, J. N.; Lee, S.; and Narayanan, S. S. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42(4):335.
- Chen, M.; Wang, S.; Liang, P. P.; Baltrušaitis, T.; Zadeh, A.; and Morency, L.-P. 2017. Multimodal sentiment analysis with word-level fusion and reinforcement learning. In *ICMI*, 163–171. ACM.
- Conneau, A.; Kiela, D.; Schwenk, H.; Barrault, L.; and Bordes, A. 2017. Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.
- Degottex, G.; Kane, J.; Drugman, T.; Raitio, T.; and Scherer, S. 2014. Covarep—a collaborative voice analysis repository for speech technologies. In *ICASSP*, 960–964. IEEE.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Gers, F. A.; Schmidhuber, J.; and Cummins, F. 1999. Learning to forget: Continual prediction with lstm. *ICANN*.
- Haq, S., and Jackson, P. J. 2011. Multimodal emotion recognition. In *MAPAS*. IGI Global. 398–423.
- Hazarika, D.; Poria, S.; Gorantla, S.; Cambria, E.; Zimmermann, R.; and Mihalcea, R. 2018. Cascade: Contextual sarcasm detection in online discussion forums. *arXiv preprint arXiv:1805.06413*.
- Hotelling, H. 1936. Relations between two set of variables. *Biometrika* 28(3-4):321–377.
- iMotions. 2017. Facial expression analysis. *imotions.com*.
- Lawrence, S.; Giles, C. L.; Tsoi, A. C.; and Back, A. D. 1997. Face recognition: A convolutional neural-network approach. *IEEE transactions on neural networks* 8(1):98–113.
- Lee, G.; Singanamalli, A.; Wang, H.; Feldman, M. D.; Master, S. R.; Shih, N. N.; Spangler, E.; Rebbeck, T.; Tomaszewski, J. E.; and Madabhushi, A. 2014. Supervised multi-view canonical correlation analysis (smvcca): Integrating histologic and proteomic features for predicting recurrent prostate cancer. *IEEE transactions on medical imaging* 34(1):284–297.
- Liang, P. P.; Liu, Z.; Zadeh, A.; and Morency, L.-P. 2018. Multimodal language analysis with recurrent multistage fusion. *arXiv preprint arXiv:1808.03920*.
- Liu, Z.; Shen, Y.; Lakshminarasimhan, V. B.; Liang, P. P.; Zadeh, A.; and Morency, L.-P. 2018. Efficient low-rank multimodal fusion with modality-specific factors. *ACL*.
- Liu, F.; Guan, B.; Zhou, Z.; Samsonov, A.; G Rosas, H.; Lian, K.; Sharma, R.; Kanarek, A.; Kim, J.; Guermazi, A.; and Kijowski, R. 2019. Fully-automated diagnosis of anterior cruciate ligament tears on knee mr images using deep learning. *Radiology* 1.
- Martin, N., and Maes, H. 1979. *Multivariate analysis*. Academic press London.
- Morency, L.-P.; Mihalcea, R.; and Doshi, P. 2011. Towards Multimodal Sentiment Analysis: Harvesting Opinions from The Web. In *ICMI 2011*.
- Poria, S.; Chaturvedi, I.; Cambria, E.; and Hussain, A. 2016. Convolutional mkl based multimodal emotion recognition and sentiment analysis. In *ICDM*, 439–448. IEEE.
- Rotman, G.; Vulić, I.; and Reichart, R. 2018. Bridging languages through images with deep partial canonical correlation analysis. In *ACL (Volume 1: Long Papers)*, 910–921.
- Sarma, P. K.; Liang, Y.; and Sethares, W. A. 2018. Domain adapted word embeddings for improved sentiment classification. *arXiv preprint arXiv:1805.04576*.
- Soleymani, M.; Garcia, D.; Jou, B.; Schuller, B.; Chang, S.-F.; and Pantic, M. 2017. A survey of multimodal sentiment analysis. *Image and Vision Computing* 65:3–14.
- Sun, Z.; Sarma, P. K.; Sethares, W.; and Bucy, E. P. 2019. Multimodal sentiment analysis using deep canonical correlation analysis. *arXiv preprint arXiv:1907.08696*.
- Tenenhaus, A., and Tenenhaus, M. 2011. Regularized generalized canonical correlation analysis. *Psychometrika* 76(2):257.
- Tsai, Y.-H. H.; Liang, P. P.; Zadeh, A.; Morency, L.-P.; and Salakhutdinov, R. 2018. Learning factorized multimodal representations. *arXiv preprint arXiv:1806.06176*.
- Tsai, Y.-H. H.; Bai, S.; Liang, P. P.; Kolter, J. Z.; Morency, L.-P.; and Salakhutdinov, R. 2019. Multimodal transformer for unaligned multimodal language sequences. *arXiv preprint arXiv:1906.00295*.
- Wang, Y.; Shen, Y.; Liu, Z.; Liang, P. P.; Zadeh, A.; and Morency, L.-P. 2018. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. *arXiv preprint arXiv:1811.09362*.
- Wang, L.; Li, Y.; and Lazebnik, S. 2016. Learning deep structure-preserving image-text embeddings. In *CVPR*, 5005–5013.
- Zadeh, A.; Zellers, R.; Pincus, E.; and Morency, L.-P. 2016. Mosi: multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos. *arXiv preprint arXiv:1606.06259*.
- Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; and Morency, L.-P. 2017. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*.
- Zadeh, A.; Liang, P. P.; Mazumder, N.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018a. Memory fusion network for multi-view sequential learning. In *AAAI*.
- Zadeh, A.; Liang, P. P.; Poria, S.; Vij, P.; Cambria, E.; and Morency, L.-P. 2018b. Multi-attention recurrent network for human communication comprehension. In *AAAI*.
- Zadeh, A. B.; Liang, P. P.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018c. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *ACL*, 2236–2246.