

Quantum Cognitively Motivated Decision Fusion Framework for Video Sentiment Analysis

Abstract

Video sentiment analysis as a decision-making process is inherently complex, involving fusion of decisions from multiple modalities and the so-caused cognitive biases. Inspired by recent advances in the emerging field of quantum cognition, we show that the sentiment judgment from one modality could be incompatible with the judgment from another, i.e., the order matters and they cannot be jointly measured to produce a final judgment. Thus the cognitive process in video sentiment analysis exhibits “quantum-like” biases that cannot be captured by classic probability theories. Accordingly, we propose a fundamentally new, quantum cognitively-motivated fusion framework for predicting sentiment judgments. In particular, we formulate utterances as quantum superposition states of positive and negative sentiment judgments, and uni-modal classifiers as mutually incompatible observables, on a complex-valued Hilbert space with positive-operator valued measures. Experiments on two benchmarking datasets illustrate that our framework significantly outperforms various existing decision level and a range of state-of-the-art content-level fusion approaches. The results also show that the concept of incompatibility allows effective handling of all combination patterns including those extreme cases that are wrongly predicted by all uni-modal classifiers.

Introduction

Video sentiment analysis is an emerging interdisciplinary area, bringing together artificial intelligence (AI) and cognitive science. It studies a speaker’s sentiment expressed by distinct modalities, i.e., linguistic, visual, and acoustic. At its core, effective modality fusion strategies are in place. Existing neural structures achieve the state-of-the-art (SOTA) (Tsai et al. 2019; Wang et al. 2019b; Dumpala et al. 2019; Zadeh et al. 2017) by integrating features after being extracted, called *feature-level* fusion. Other approaches simulate logic reasoning and human cognitive biases (Morvant, Habrard, and Ayache 2014; Glodek et al. 2011b,a) by aggregating decisions of uni-modal classifiers into a joint decision, called *decision-level* fusion. Additionally, hybrid fusion approaches benefit from the advantages of both strategies. In this paper, we target at the generally less effective but more flexible decision-level fusion.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Video sentiment analysis is inherently a complex human cognition process. Recent research in cognitive science found that in some cases human decision making could be highly irrational (Tversky and Kahneman 1983), and such behaviour does not always obey the classical (Kolmogorov) axioms of probability (Kolmogorov 1950) and utility theory (Morgenstern 1949). On the other hand, preliminary work show that the mathematical formalism of Quantum Mechanics (QM) can successfully address paradoxes of classical probability theory in modelling human cognition (Busemeyer and Bruza 2012). Conceptually, quantum cognition challenges the notion that user’s cognitive states underpinning the decisions have pre-defined values and that a measurement merely records them. Instead, the cognitive system is fundamentally uncertain and in an indefinite state. The act of measurement would then create a definite state out of the indefinite state and changes the state of the system.

We hypothesise that uni-modal sentiment judgments do not happen independently, like a pre-defined value being read out of the internal cognitive state. They are rather constructed at the point of information interaction and thus influenced by the other modalities, which serve as a context for the inference of sentiment judgment for the current modality. For example, there might be cases that the order of different decision perspectives, e.g., when someone focuses first on the linguistic and then on the visual perspective, or vice versa, can lead to controversial sentiment judgements. That is, the measurement from the first perspective provides a context that affects the subsequent one, influencing the probabilities used to compute the utility function of multimodal sentiment decision. In this case, we say that these two decision perspectives are *incompatible* with each other. That implies that judgements over different modalities cannot be measured jointly, and quantum probability should hence be in place (Uprety et al. 2020). We argue that video sentiment analysis could benefit from the generalized framework of quantum cognition by capturing the cases of incompatibility that are not possible with classical probability.

In this work, we introduce a novel decision-level fusion framework inspired by quantum cognition (Fell et al. 2019). The goal is to predict the sentiment of utterances in videos, associated with linguistic, visual, and acoustic streams. In particular, we formulate an utterance as a *quantum superposition* state of positive and negative sentiments (i.e., it can be

positive and negative at the same time until it is judged under a specific context), and uni-modal classifiers as *mutually incompatible observables*, on a shared complex-valued Hilbert space \mathcal{H} spanned by distinct uni-modal sentiment bases. In this work, we take advantage of incompatibility to influence the uni-modal decisions, when they are under high uncertainty, to finally infer multimodal sentiment judgments. It is important to note that the framework produces a generalization form of classical probabilities, allowing for both compatible and incompatible sentiment decisions.

Specifically, we resolve the incompatibility issue through Positive-Operator Valued Measures (POVMs) to approximate the sentiment of uni-modal classifiers simultaneously. In practice, we estimate the complex Hilbert Space and uni-modal observables from training data, and then establish the final multimodal sentiment state of a test utterance from the learned uni-modal observables.

To our best knowledge, this is the first quantum cognitively inspired theoretical approach, with practical implementations, that investigates and models the incompatibility of sentiment judgments for video sentiment analysis.

Extensive evaluation on two widely used benchmarking datasets, namely CMU-MOSI(Zadeh et al. 2016) and CMU-MOSEI(Bagher Zadeh et al. 2018), show that our framework significantly outperforms not only various representative decision-level fusion baselines, but also a range of SOTA content-level fusion approaches for video sentiment analysis. We also show that the framework is able to make correct sentiment judgments even for the cases where all uni-modal classifiers give wrong predictions.

Related Work

Video Sentiment Analysis

Existing recurrent neural structures have achieved SOTA results for utterance-level video sentiment analysis. One example is memory networks incorporating multimodal hidden units of preceding timestamps with the inputs (Wang et al. 2019b; Liang et al. 2018; Zadeh et al. 2018). Tensor-based operations had also been exploited to entangle (Zadeh et al. 2017) and disentangle (Liu et al. 2018; Barezi and Fung 2019) multimodal information. Inspired from successful trends in NLP, some approaches introduced fuzzy logic (Chaturvedi et al. 2019) and encoder-decoder structures in sequence-to-sequence learning, translating a target modality to a source modality (Pham et al. 2019; Dumpala et al. 2019). Multimodal transformer (Tsai et al. 2019) achieved the SOTA for video sentiment analysis. It is worth mentioning that there exist research focusing on unaligned time-series streams (Wang et al. 2019b), but those approaches are beyond the scope of this work.

Quantum-inspired Representation Learning

The application of quantum theory in representation learning began after van Rijsbergen’s pioneering work (Van Rijsbergen 2004) by integrating geometric spaces and probabilistic spaces into a unified theoretical framework. The mathematical formalism of QM has been exploited for Information Retrieval (IR), NLP, and multimodal representa-

tion learning tasks. Among them, quantum formalism has been successfully utilised for modelling word dependencies through density matrices (Sordani, Nie, and Bengio 2013) and formulating the semantic composition of words (Sordani, He, and Nie 2013) in IR tasks. Quantum-inspired models have also been introduced to address NLP tasks. Early work deployed neural networks to learn *entangled* text representations for question-answering tasks (Zhang et al. 2018a). The exploitation of quantum concepts in complex-valued Hilbert spaces led to improved performance and better interpretability (Wang et al. 2019a; Li, Wang, and Melucci 2019). A quantum-inspired recurrent network addressed interactions across speakers in conversations (Zhang et al. 2019). Quantum-inspired strategies have been utilised for multimodal representation learning as well: Wang et al. (Wang, Song, and Kaliciak 2010) proposed a tensor-based representation to retrieve image-text documents. Moreover, there have been studies that investigated *quantum interference* (Zhang et al. 2018b) and *non-classical correlations* (Gkoumas, Uprety, and Song 2018) to address the decision-level modality fusion. Recently, a quantum-inspired neural framework achieved the SOTA performance for utterance-level video sentiment analysis (Li et al. 2020). Unlike these existing works, in this paper, we propose a quantum cognition inspired theoretical framework capturing cognitive biases that are not possible with classical probability, through the concept of incompatibility.

Background of Quantum Cognition

In this section, we introduce the key concepts from quantum cognition (Busemeyer and Bruza 2012; Fell et al. 2019), which we exploited to construct the proposed framework.

Hilbert Space

Quantum cognition exploits an infinite complex-valued vector space, called Hilbert space \mathcal{H} , in which the state of a quantum system is represented as a vector. Different from classical probability, quantum probability events are defined as orthonormal basis states. A projective geometric structure establishes relationships between states vectors and basis states (Melucci 2008). The same Hilbert space can be represented by different sets of orthonormal basis states, and the same state can be defined over different sets of orthonormal basis states.

In consistency with QM, we adopt the widely-used *Dirac Notations* for the mathematical formalism of quantum cognition. A complex-valued *unit* vector \vec{u} and its conjugate transpose \vec{u}^{*T} are denoted as a *ket* $|u\rangle$ and a *bra* $\langle u|$, respectively. The inner product of two vectors $|u\rangle$ and $|v\rangle$ is defined by $\langle u|v\rangle$, while $|u\rangle\langle u|$ and $|v\rangle\langle v|$ define operators.

Quantum Superposition

Quantum superposition is one of the fundamental concepts in QM, which describes the uncertainty of a single particle. In the micro world, a particle like a photon can be in multiple mutually exclusive basis states simultaneously with a probability distribution. A general pure state $|\psi\rangle$ is a vector

on the unit sphere, represented by

$$|\psi\rangle = w_1 |e_1\rangle + \dots + w_n |e_n\rangle. \quad (1)$$

Where $|e_1\rangle, \dots, |e_n\rangle$ are *basis states* that form an orthogonal basis of the Hilbert Space, and the *probability amplitudes* $\{w_j\}$ are complex scalars with $\sum_j |w_j|^2 = 1$, with $|\cdot|$ being the modulus of a complex number. $|\psi\rangle$ is a *superposition state* when it is not identical to a certain basis state $|e_j\rangle$. In particular, in a two-dimensional Hilbert Space \mathcal{H}_2 spanned by basis states $|0\rangle$ and $|1\rangle$, a pure state $|\psi\rangle$ is represented as

$$|\psi\rangle = \cos \frac{\theta}{2} |0\rangle + e^{i\phi} \sin \frac{\theta}{2} |1\rangle \quad (2)$$

while $\theta, \phi \in [0, 2\pi]$ and i is the imaginary number satisfying $i^2 = -1$. Eq. 2 uniquely expresses any pure state on \mathcal{H}_2 .

Measurement

Measurement is another fundamental concept in quantum cognition for calculating quantum probabilities. In QM, Projection-Valued Measure (PVM) removes a system state from uncertainty to a precise event, by projecting a state to its corresponding certain basis state¹. In the absence of measurement, there is uncertainty in the state in that it takes all possible measurement values simultaneously. After measurement, the state *collapses* onto precisely one possible value. However, PVMs on subsystems of larger systems cannot be described by a PVM acting on the subsystem itself. Positive-Operator Valued Measure (POVM) overcomes this constraint, by associating a positive probability for each measurement outcome, ignoring the post-measure state (Nielsen and Chuang 2011). That is, POVM is a generalization of PVM, providing mixed information of a state for the entire ensemble of subsystems.

Mathematically speaking, a POVM M is a set of *Hermitian positive semi-definite* operators $\{E_i\}$ on a Hilbert space \mathcal{H} that sum to the identity operator, i.e., $\sum_i E_i = \mathbb{1}$. For a generic pure state's density matrix ρ , where $\rho = |\psi\rangle\langle\psi|$, the probability with respect to E_i is computed as

$$P(i) = \text{Tr}(E_i \rho) = \langle\psi| E_i |\psi\rangle, \quad (3)$$

and $\sum P(i) = 1$.

In the case of measuring a state on a two-dimensional Hilbert Space \mathcal{H}_2 (see Eq. 2), the POVM is associated with the following operators (Busch 1986):

$$E_+ = \frac{\eta}{2} \mathbb{1} + (1 - \eta) |1\rangle\langle 1| \quad (4)$$

$$E_- = \frac{\eta}{2} \mathbb{1} + (1 - \eta) |0\rangle\langle 0|, \quad (5)$$

where $\mathbb{1}$ and η stand for the identity matrix and noise parameter respectively. The value $\eta \in [0, 1]$ determines the probability that the measurement fails due to the system-apparatus correlation or incompatibility (Liang, Spekkens, and Wiseman 2011). When $\eta = 0$, the measurement apparatus exerts

¹For simplicity, we choose a nomenclature instead of a strict definition. Please refer to (Nielsen and Chuang 2011) for a strict definition of projection measurement.

no influence on the measurement, and we have an approximate measurement. When $\eta = 1$, the output of measurement is completely random.

Incompatibility

The concept of *incompatibility* is applicable in a Hilbert space only. Each basis state, defining an event, has a projector Π to evaluate the event. In contrast to classical probability, the conjunction of two events is not necessarily commutative (Busemeyer and Wang 2018). Suppose Π_A and Π_B are two sequential measurements for A and B events. In quantum cognition, the joint probability distribution of two events equals the product of the two projectors Π_A and Π_B , corresponding to the basis state $A \cap B$. If $\Pi_A \Pi_B = \Pi_B \Pi_A$, then the two events are called *compatible*. However, if $\Pi_A \Pi_B \neq \Pi_B \Pi_A$, then their product is not a projector, and the two events do not commute, that is, they are *incompatible*. Incompatibility implies that the two measurements cannot be accessed jointly without disturbing each other. Classical probability can not capture such disturbance, assuming that measurements are always compatible and commute. However, the mathematical formalism of quantum probability allows for both compatible and incompatible measurements (Hughes 1989; Uprety et al. 2020). Thus, it is a generalization of classical probability theory.

Quantum Cognition-inspired Framework

We now introduce the proposed quantum cognition-inspired theoretical framework for video sentiment analysis.

Task Formulation

Due to space limitation, this work targets at the binary video sentiment analysis task. However, the framework is extendable, by adopting an one-vs-all multiclass classification strategy. Each utterance $U_i \in \{U_1, \dots, U_N\}$ is associated with linguistic, visual and acoustic features $U_i = \{X_{i,l}, X_{i,v}, X_{i,a}\}$ and a positive or negative sentiment label $y_i \in [-1, 1]$. The objective is to establish a function, mapping an utterance U_i to its corresponding sentiment label.

Sentiment Hilbert Space

The framework is defined on a *Sentimental Hilbert Space* \mathcal{H}_{senti} , which is a 2-dimensional vector space spanned by the basis states $\{|+\rangle, |-\rangle\}$. The basis states $|+\rangle, |-\rangle$ correspond to positive and negative sentiments, respectively. We represent an utterance U_k as a pure state $|S_{U_k}\rangle$ (in short $|S\rangle$) on \mathcal{H}_{senti} . The uni-modal sentiment classifiers (denoted as L, V, A) are formulated as mutually incompatible observables (see Figure 1). The utterance can be represented under different sets of basis states, i.e., uni-modal (L, V, A) and multimodal (F) basis states in Figure 1. The observables are not orthogonal since modalities are not independent, but highly correlated.

Utterance Representation

An utterance is represented as a pure state $|S\rangle$ of positive and negative sentiments on \mathcal{H}_{senti} (see Figure 1). On this 2-dimensional Hilbert space, it adopts a generic representation

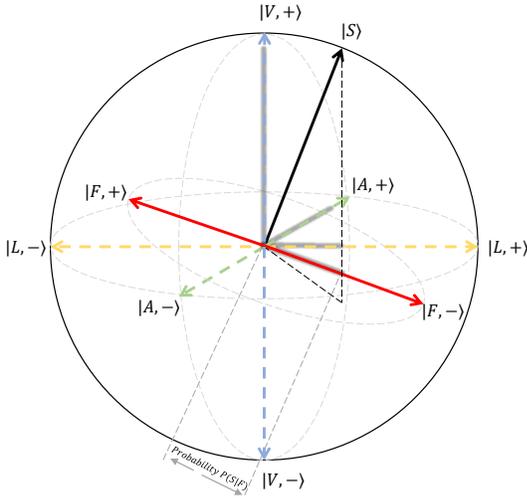


Figure 1: The Sentimental Hilbert Space. An utterance is represented as a pure state $|S\rangle$ belonging to the surface of a unit sphere. The two opposed unit vectors represent positive and sentiment judgments. The associated uni-modal sentiment observables $\hat{L}, \hat{V}, \hat{A}$ and tri-modal observable \hat{F} are mutually incompatible. Shadowed basis vectors imply projections of $|S\rangle$ on the corresponding bases, i.e., probability of events.

$$|S\rangle = \cos \frac{\theta_S}{2} |+\rangle + e^{i\phi_S} \sin \frac{\theta_S}{2} |-\rangle, \quad (6)$$

where $\theta_S, \phi_S \in [0, 2\pi]$. According to the Born's rule (Born 1926), the probability of an utterance being positive and negative is $P(+)=|\langle S|+\rangle|^2=\cos^2\frac{\theta_S}{2}$ and $P(-)=|\langle S|-\rangle|^2=\sin^2\frac{\theta_S}{2}$, with $\cos^2\frac{\theta_S}{2}+\sin^2\frac{\theta_S}{2}=1$. The *relative phase* ϕ_S gives rise to different results from classical case. Crucially, relative phases play an important role in correlations between incompatible observables.

Sentiment Decisions

We formulate uni-modal sentiment decisions as mutually incompatible observables on $\mathcal{H}_{\text{sent}}$, namely \hat{L}, \hat{V} , and \hat{A} for linguistic, visual and acoustic modalities respectively (see Figure 1). For the binary sentiment analysis task, each observable is associated with two eigenstates and two eigenvalues, with common eigenvalues of 1 and -1 for positive and negative sentiments. In that case, incompatibility falls under different sets of eigenstates $\{|M, +\rangle, |M, -\rangle\}$ defining a uni-modal basis, where modality $M \in \{L, V, A\}$. Following Eq. 6, we express the eigenstates as

$$|M, +\rangle = \cos \frac{\theta_M}{2} |+\rangle + e^{i\phi_M} \sin \frac{\theta_M}{2} |-\rangle \quad (7)$$

$$|M, -\rangle = \sin \frac{\theta_M}{2} |+\rangle - e^{i\phi_M} \cos \frac{\theta_M}{2} |-\rangle \quad (8)$$

with $\theta_M, \phi_M \in [0, 2\pi]$. The eigenstates form an orthonormal basis, with $\langle M, +|M, +\rangle = \langle M, -|M, -\rangle = 1$ and $\langle M, +|M, -\rangle = \langle M, -|M, +\rangle = 0$.

A general observable \hat{O} can be decomposed to its eigenstates $\{|\lambda_i\rangle\}$ of the orthonormal basis as $\hat{O} = \sum \lambda_i |\lambda_i\rangle \langle \lambda_i|$, where eigenvalues $\{\lambda_i\}$ are possible values that a state can take for the corresponding events after measurement. Thus, the uni-modal observables are defined as follows:

$$\hat{M} = (+1) |M, +\rangle \langle M, +| + (-1) |M, -\rangle \langle M, -| \quad (9)$$

where observable $\hat{M} \in \{\hat{L}, \hat{V}, \hat{A}\}$. Similarly, the observable for the final sentiment decision \hat{F} is

$$\hat{F} = (+1) |+\rangle \langle +| + (-1) |-\rangle \langle -| \quad (10)$$

which spans the $\mathcal{H}_{\text{sent}}$ and is incompatible with all uni-modal observables.

Following the projective geometric structure, the result probability of a measurement on an eigenstate equals the projection of the state onto it, that is, the squared inner product of the vectors, e.g., $|\langle S|M, +\rangle|^2$ for unimodal positive sentiment and $|\langle S|+\rangle|^2$ for final (multimodal) positive sentiment. The measurement probabilities under \hat{L} stand for its sentiment under linguistic modality, as so forth for the other modalities. Finally, its multimodal sentiment is determined by the observable \hat{F} . In Figure 1, the sentiment judgment is positive in terms of uni-modal observables (projections are visualized as shadows in Figure 1), yet is negative in terms of the multimodal observable due to incompatibility.

Method

This section presents a methodology that operationalizes the proposed framework. Traditionally, in the physical sciences, the study of mathematical problems involves modelling methods leveraging a combination of approximation techniques. In this work, we exploit statistical information from the data to learn the sentimental Hilbert Space described in the previous section, so as to leverage the incompatible observables to determine the sentiment of an utterance. Overall, the pipeline consists of three steps: (1) we first estimate the generic uni-modal observables \hat{M} from the training data; (2) then we construct the sentiment state for each test utterance $|S_T\rangle$ from the learned uni-modal observables and uni-modal sentiment prediction results; (3) finally, we judge the sentiment from the multimodal observable \hat{F} . In the remaining part of the section, we elaborate the methodology for each step.

Observable Estimation

The uni-modal observables are constructed from the overall statistics of the training data. These values are mapped to their quantum expressions to estimate the parameters of the uni-modal observables. In particular, the uni-modal observables and pure state should submit to the following properties: I) the pure state should conform to the statistics of the dataset, II) the uni-modal sentiment measurement results should be consistent with the ratio of positive and negative samples in the training subsets, and III) quantum correlations between observables should be aligned to classical correlations of the per-sample prediction results, derived from the training data.

To facilitate the construction of uni-modal observables, we introduce a pure state as follows:

$$|G\rangle = \cos \frac{\theta_G}{2} |+\rangle + e^{i\phi_G} \sin \frac{\theta_G}{2} |-\rangle, \quad (11)$$

which describes the extent to which the dataset is unbalanced for positive and negative labels. By Born's rule (Born 1926), the probability of positive judgment is:

$$P(+)=|\langle +|G\rangle|^2\approx\frac{\#pos}{N}, \quad (12)$$

where $\#pos$ is the number of true positive utterances in the training set and N the size of the training set. Eq.12 implies

$$\cos^2 \frac{\theta_G}{2} \approx \frac{\#pos}{N}, \quad (13)$$

since the quantum probability equals the squared amplitude of a state (see Section Background or (Melucci 2015)).

According to the second property, the probability of a positive sentiment judgment for each modality is given by

$$P_M(+)=|\langle M,+|G\rangle|^2\approx\frac{\#M_{pos}}{N}, \quad (14)$$

where $\#L_{pos}$, $\#V_{pos}$ and $\#A_{pos}$ equals the number of true positive utterances for each modality in the training set. Combining Eq.7, Eq.11 and Eq.14, the probability of the positive sentiment judgment for each modality is

$$\begin{aligned} &\cos^2 \frac{\theta_M}{2} \cos^2 \frac{\theta_G}{2} + \sin^2 \frac{\theta_M}{2} \sin^2 \frac{\theta_G}{2} \\ &+ \frac{1}{2} \sin \theta_M \sin \theta_G \cos(\phi_M - \phi_G) \approx \frac{\#m_{pos}}{N}. \end{aligned} \quad (15)$$

Finally, we looked into the correlations between pairs of uni-modal observables, where the relative phases play an important role. From a quantum statistics point of view, the correlation of observables for two modalities M_1, M_2 is given by $(|\langle M_1,+|M_2,+ \rangle|^2 + |\langle M_1,-|M_2,- \rangle|^2 - |\langle M_1,+|M_2,- \rangle|^2 - |\langle M_1,-|M_2,+ \rangle|^2)$. A more detailed explanation about the derivation is included in Appendix I. It should in principle be aligned to the classical correlations derived from the data. Hence we have

$$\begin{aligned} &\frac{1}{2} (|\langle M_1,+|M_2,+ \rangle|^2 + |\langle M_1,-|M_2,- \rangle|^2 \\ &- |\langle M_1,+|M_2,- \rangle|^2 - |\langle M_1,-|M_2,+ \rangle|^2) \\ &\approx corr(M_1, M_2), \end{aligned} \quad (16)$$

where $M_1 \neq M_2 \in \{L, V, A\}$ and $corr(M_1, M_2)$ is a classical correlation of the per-sample prediction results based on modalities M_1 and M_2 , which is computed from the training data. When M_1 and M_2 give exactly same predictions, the correlation $corr(M_1, M_2) = 1$. Therefore, $|\langle M_1,+|M_2,+ \rangle| = |\langle M_1,-|M_2,- \rangle|^2 = 1$ and $|\langle M_1,+|M_2,- \rangle|^2 = |\langle M_1,-|M_2,+ \rangle|^2 = 0$, so the value 1 is also produced from the quantum side. Similarly, a value of -1 is obtained for both sides when the two modalities give totally opposite predictions, indicating the maximum negative correlation. Hence, Eq.16 gives three equations for distinct pairs of modalities. For example, for linguistic-visual correlation, Eq.16 results in

$$\begin{aligned} &\cos \theta_L \cos \theta_V + \sin \theta_L \sin \theta_V \cos(\phi_L - \phi_V) \\ &\approx corr(L, V), \end{aligned} \quad (17)$$

as so forth for the $\{L, A\}, \{A, V\}$ modality pairs.

To wrap up, taking into account the number of positive sentiments in the training set and correlations across different pairs of modalities, we get 7 equations from Eq. 13, Eq. 15 and Eq. 16, and 8 unknown variables $\{\theta_G, \theta_L, \theta_V, \theta_A, \phi_G, \phi_L, \phi_V, \phi_A\}$. As all equations rely only on the differences between the relative phases rather than their absolute values, we can safely set $\phi_G = 0$ without loss of information. Hence a unique solution of $\{\theta_G, \theta_L, \theta_V, \theta_A, \phi_L, \phi_V, \phi_A\}$ can be produced, and accordingly determining the uni-modal observables \hat{L}, \hat{V} , and \hat{A} .

Utterance State Estimation

After having uni-modal observables calculated as described above, we need to estimate the state for each test utterance. For a specific test utterance denoted as

$$|S_T\rangle = \cos \frac{\theta_T}{2} |+\rangle + e^{i\phi_T} \sin \frac{\theta_T}{2} |-\rangle, \quad (18)$$

the uni-modal predictions can be exploited to estimate the values of θ_T, ϕ_T . However, since the observables \hat{L}, \hat{V} , and \hat{A} are mutually incompatible, the measurements results cannot be accessed simultaneously. To that end, we propose utilize POVMs to get the results of all incompatible measurements simultaneously (Uola et al. 2016). In particular, we construct sample-specific POVMs for each uni-modal measurement, applying unsharp (weak) projections (Busch 1986) without disturbing the observables. Finally, the operators are constructed as follows:

$$E_{\pm}^M = \frac{\eta_T}{2} \mathbb{I} + (1 - \eta_T) |M, \pm\rangle \langle M, \pm|, \quad (19)$$

where $\eta_T \in [0, 1]$ is specific to sample t , since each utterance interacts with the apparatus in a different manner. We apply uni-modal POVMs on the test utterance to measure the sentiment of utterance in terms of each modality, that is,

$$\langle S_T | E_{\pm}^M | S_T \rangle \approx P_{T,M}(\pm), \quad (20)$$

where $P_{T,M}(\pm)$ are uni-modal probabilities for the positive sentiment judgment. Eq.20 gives a system with three equations, each equation for a distinct modality, and three unknown variables $\{\theta_T, \phi_T, \eta_T\}$. Solving the system would allow us to construct the state $|S_T\rangle$.

Multimodal Sentiment Measurement

The sentiment of a test utterance $|S_T\rangle$ is measured by Eq. 10. The results are $P_T(+)=\cos^2 \frac{\theta_T}{2}$ and $P_T(-)=\sin^2 \frac{\theta_T}{2}$. The sentiment of S_T is hence positive if $\cos^2 \frac{\theta_T}{2} > 0.5$ and negative otherwise.

Experiments

We evaluate the proposed framework on two benchmarking datasets, namely, CMU Multimodal Opinion-level Sentiment Intensity (CMU-MOSI) (Zadeh et al. 2016) and CMU Multimodal Opinion Sentiment and Emotion Intensity (CMU-MOSEI) (Bagher Zadeh et al. 2018). For both datasets, each sample is labelled by three human annotators with a 7-level ratio score from -3 (highly negative) to 3

(highly positive). In this work, we adopt the binary accuracy (i.e., Acc_2 : positive sentiment if the human annotation score ≥ 0 , and negative sentiment if the human annotation score < 0), and $F1$ score. The statistical details of the datasets and feature extraction are included in Appendix II and III.

Baselines

We compared with robust approaches on both decision-level and feature-level modality fusion approaches.

Decision-level: We first trained neural uni-modal classifiers. In particular, we used Bi-GRU layers (Cho et al. 2014) with forward and backward state concatenation, followed by fully connected layers. The outputs gave linguistic, visual, and acoustic embeddings $\{L, V, A\} \in \mathbf{R}^d$, where d was the number of neurons in dense layers. Then, self-attentions were computed for each uni-modal dense representation, by calculating the scaled dot-product (Vaswani et al. 2017). Finally, each attentive uni-modal representation was fed into two fully connected layers, followed by a softmax layer to obtain sentiment judgments. We leave the final settings in Appendix IV. The unimodal results are then fed into the multimodal meta-fusion approaches. We compare with a range of baseline fusion approaches:

- **Voting** was used to aggregate the outputs of the uni-modal classifiers. In particular, we applied a) *Hard Voting*, via majority voting, b) *Weighted Majority Voting*, by assigning weights to each uni-modal classifier and taking their average, and c) *Soft Voting*, by averaging the predicted probabilities, to infer multi-modal sentiment judgments.
- **Single models** exploited supervised machine learning algorithms as meta fusion approaches of the uni-modal classifiers. For both tasks, we chose the most effective models, namely, a) *Logistic Regression*, b) *Support Vector Machine (SVM)*, and c) *Gaussian Naive Bayes (GaussianNB)*, from a pool of supervised learning algorithms.
- **Ensemble methods** combined learning algorithms, selecting the optimum combination from a pool of models. We explored stacking, backing, and boosting strategies (Ponti Jr 2011); a) for stacking, single models were stacked together and the hard voted method computed predictions, b) for bagging, a number of estimators were aggregated by majority voting, and c) for boosting, we applied AdaBoost classifier (Freund and Schapire 1997).
- **A Deep Fusion** approach combined the confidence scores of uni-modal classifiers along with the complementary scores as inputs to a deep neural network, followed by a sigmoid layer, which made the final prediction (Nojavanasghari et al. 2016).

Feature-level: We also compared the framework with a range of SOTA feature-level fusion approaches.

- For **SOTA**, we replicated a) *MULT* (Tsai et al. 2019), consisting of pairwise crossmodal transformers, the outputs of which are concatenated to build the multimodal embedding utterance, b) *RAVEN* (Wang et al. 2019b), an RNN based framework with an attention gating mechanism to model crossmodal interactions, and c) *TFN* (Zadeh

et al. 2017), a tensor-based neural network that a multi-dimensional tensor captures uni-modal, bi-modal, and tri-modal interactions across distinct modalities.

- **QMF** (Li et al. 2020) is a complex-valued neural framework, which represents utterances as superposed states, and incorporates modalities through a tensor operator.

Experiment Settings

We conducted the experiments on the same uni-modal classifiers trained for the decision-level baseline approaches. We estimated the uni-modal observables from training plus validation sets, and then we used the learnt observables for predicting the utterance sentiment on the test set. We used Pearson correlation for modelling classical correlations. In case the equation systems of the framework did not have solutions, the MATLAB *fsolve* function was used to generate a numerical solution. In particular, we randomly initialized the parameters $\{\theta_G, \theta_L, \theta_V, \theta_A, \phi_L, \phi_V, \phi_A\} \in [0, 2\pi]$ for uni-modal observable estimation, and $\{\theta_T, \phi_T\} \in [0, 2\pi], \eta_T \in [0, 1]$ for utterance state estimation. The random initialization was repeated for 200 times to obtain the optimum solutions by calculating the minimum sum of squared loss.

Comparative Analysis of Results

For both tasks, we present the comparison results between the proposed framework and various decision-level fusion strategies in Table 1. For CMU-MOSEI, all approaches attained an improved performance as compared to the performance of CMU-MOSI task. We suspect this is because MOSEI is a much larger dataset. Overall, Weighted Voting is the best-performing approach among the voting-based aggregations, Logistic Regression among the supervised learning algorithms, and Stacking among the ensemble learning methods. For both tasks, Stacking and Bagging are the most effective baseline decision-level fusion strategies. For CMU-MOSI, the proposed framework attain an increased accuracy of 84.6% as compared to 78.4% of Stacking. That is, a significant improvement of 6.2% ($p - value < 0.05$). For CMU-MOSEI, the framework reaches an increased accuracy of 84.9% as compared to 82.2% of Stacking, i.e., a significant improvement of 2.7% ($p - value < 0.05$).

Approach	CMU-MOSI		CMU-MOSEI	
	Acc_2	$F1$	Acc_2	$F1$
Hard Voting	67.5	65.4	71.5	83.3
Weighted Voting	74.6	71.6	81.3	87.8
Soft Voting	75.2	71.9	77.5	86.2
SVM	77.4	72.9	81.7	87.9
Logistic Regression	78.0	73.8	81.9	88.0
GaussianNB	76.7	71.6	80.9	86.8
Stacking	78.4	75.1	82.2	88.1
Bagging	78.1	73.6	82.0	88.0
Boosting	77.7	74.0	81.7	87.7
Deep Fusion	77.8	77.7	81.9	81.3
Proposed Framework	84.6 ($\uparrow 6.2$)	84.5 ($\uparrow 6.8$)	84.9 ($\uparrow 2.7$)	91.1 ($\uparrow 3.0$)

Table 1: Effectiveness of decision-level fusion approaches on CMU-MOSI and CMU-MOSEI. Best results are highlighted in boldface. Numbers in parentheses indicate relative percentage improvement over the next best model.

Table 2 presents the comparison results between the introduced framework and various feature-level fusion approaches. For CMU-MOSI, TFN (Zadeh et al. 2017) is the most effective among the baseline feature-level fusion approaches. The proposed framework attain an improvement in accuracy by 3.4% (see Table 2). For CMU-MOSEI, RAVEN (Wang et al. 2019b) attains the highest accuracy among the baselines. The proposed framework yields an increased accuracy of 84.9% as compared to 80.2% of RAVEN, i.e., 4.7% improvement.

We notice that the decision-level feature strategies achieve better performance than the feature-level neural approaches on CMU-MOSEI. This implies that discriminative learning approaches can benefit from large datasets, whereas neural approaches lead to overfitting. We also observe that the proposed framework achieves a similar performance both on CMU-MOSI and CMU-MOSEI, even though CMU-MOSI is a relatively balanced dataset. That is, our framework can effectively cope with both skewed and balanced datasets.

Approach	CMU-MOSI		CMU-MOSEI	
	Acc_2	$F1$	Acc_2	$F1$
MuT (Tsai et al. 2019)	80.2	79.5	80.0	79.8
RAVEN (Wang et al. 2019b)	78.6	78.6	80.2	79.9
TFN (Zadeh et al. 2017)	81.2	80.8	77.8	77.8
QMF (Li et al. 2020)	80.7	79.7	79.7	79.6
Proposed Framework	84.6 (\uparrow 3.4)	84.5 (\uparrow 3.7)	84.9 (\uparrow 4.7)	91.1 (\uparrow 11.2)

Table 2: Effectiveness of feature-level fusion approaches.

Ablation Tests Table 3 shows the results of our ablation study. The first three rows list the performance of uni-modal classifiers when no crossmodal interactions are modelled. The linguistic modality is the most predictive due to the use of word embedding trained on large corpora. For CMU-MOSEI, the linguistic classifier outperforms all the feature-level and voting-based fusion approaches, illustrating the robustness of uni-modal classifiers.

Approach	CMU-MOSI		CMU-MOSEI	
	Acc_2	$F1$	Acc_2	$F1$
Linguistic Only	77.1	72.3	81.5	87.8
Visual Only	54.7	48.4	71.1	83.0
Acoustic Only	56.1	60.0	71.2	83.1
Proposed Framework	84.6 (\uparrow 7.5)	84.5 (\uparrow 12.2)	84.9 (\uparrow 3.4)	91.1 (\uparrow 3.3)

Table 3: Comparison with uni-modal sentiment analysis classifiers.

As a second set of ablation experiments, we test the proposed framework when only bimodal dynamics are present. We present the result in Table 4, which shows the linguistic and acoustic dynamics are the most informative. However, trimodal dynamics outperform all possible bimodal combinations, yielding an improvement of accuracy by 5.0% for CMU-MOSI, and 2.2% for CMU-MOSEI.

Effect of Incompatibility We conduct a further analysis to investigate the effectiveness of incompatibility. We first identified all the cases that are correctly predicted by one out of the eleven decision-level fusion approaches. In total,

Approach	CMU-MOSI		CMU-MOSEI	
	Acc_2	$F1$	Acc_2	$F1$
Framework _{Linguistic, Visual}	78.2	74.3	82.1	88.4
Framework _{Linguistic, Acoustic}	79.6	75.1	82.7	89.2
Framework _{Visual, Acoustic}	55.1	55.2	70.8	82.7
Proposed Framework	84.6 (\uparrow 5.0)	84.5 (\uparrow 9.4)	84.9 (\uparrow 2.2)	91.1 (\uparrow 1.9)

Table 4: Comparison of the framework with its variants.

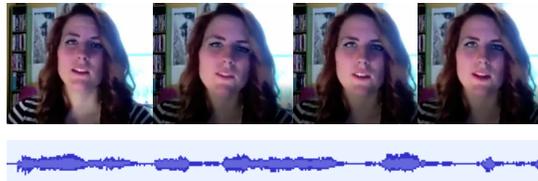


Figure 2: Visual-acoustic content of an incompatible case.

there were 33 such cases on CMU-MOSI and 1547 ones on CMU-MOSEI test sets. The proposed framework gives correct predictions for 31 cases out of 33 on CMU-MOSI and all the 1547 cases on CMU-MOSEI. Furthermore, we analyze the cases that all uni-modal classifiers give wrong sentiment judgments, but the proposed framework successfully fused them giving correct predictions. There are 39 such utterances out of 686 on the CMU-MOSI and 633 utterances out of 4643 on the CMU-MOSEI subsets.

Case Study We illustrate the visual-acoustic content of an incompatible case of the utterance “*I mean even if you don’t have kinds*” in Figure 2. The linguistic state by itself is in an indefinite state, which results in a superposition of sentiment judgments. Similarly, the visual-acoustic content is under uncertainty since the content is neutral. Indeed, all uni-modal classifiers predicted a negative sentiment judgment, inferring a probability less than 0.5, yet very close to the decision boundary of 0.5. This superposition of uni-modal beliefs, i.e., positive and negative sentiment at the same time until they are judged under a specific context, results in the occurrence of incompatibility. Under the high levels of uncertainty, incompatibility influences the uni-modal judgments and successfully predicts a positive multimodal sentiment judgment. This phenomenon is the core of the framework and the reason it achieves such high performance.

Conclusions

We introduced an effective fusion strategy inspired by quantum cognition. We formulated utterances as states and uni-modal decisions as mutually incompatible observables in a complex-valued sentimental Hilbert space. The incompatibility captures cognitive biases that are otherwise not possible with classical probability. The proposed framework has been shown able to all combination patterns including the cases where all uni-modal classifiers gave wrong sentiment judgments. Therefore, the proposed approach achieved an improved performance over SOTA feature-level and decision-level modality fusion approaches. In the future, we would like to investigate the framework on the video emotion recognition and conversational sentiment analysis tasks.

References

- Bagher Zadeh, A.; Liang, P. P.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018. 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)*, 2236–2246. Melbourne, Australia. doi: 10.18653/v1/P18-1208.
- Barezi, E. J.; and Fung, P. 2019. Modality-based Factorization for Multimodal Fusion. In *Proceedings of the 4th Workshop on Representation Learning for NLP, Repl4NLP@ACL 2019, Florence, Italy, August 2, 2019*, 260–269. Association for Computational Linguistics. doi: 10.18653/v1/w19-4331.
- Born, M. 1926. Quantenmechanik der stoßvorgänge. *Zeitschrift für Physik* 38(11-12): 803–827.
- Busch, P. 1986. Unsharp reality and joint measurements for spin observables. *Physical Review D* 33(8): 2253.
- Busemeyer, J. R.; and Bruza, P. D. 2012. *Quantum models of cognition and decision*. Cambridge University Press.
- Busemeyer, J. R.; and Wang, Z. 2018. Hilbert space multidimensional theory. *Psychological review* 125(4): 572.
- Chaturvedi, I.; Satapathy, R.; Cavallari, S.; and Cambria, E. 2019. Fuzzy commonsense reasoning for multimodal sentiment analysis. *Pattern Recognit. Lett.* 125: 264–270. doi: 10.1016/j.patrec.2019.04.024.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Dumpala, S. H.; Sheikh, I.; Chakraborty, R.; and Kopparapu, S. K. 2019. Audio-visual fusion for sentiment classification using cross-modal autoencoder. NIPS.
- Fell, L.; Dehdashti, S.; Bruza, P.; and Moreira, C. 2019. An Experimental Protocol to Derive and Validate a Quantum Model of Decision-Making. In *Proceedings of the 41th Annual Meeting of the Cognitive Science Society, CogSci 2019: Creativity + Cognition + Computation, Montreal, Canada, July 24-27, 2019*, 1724–1730. cognitivesciencesociety.org.
- Freund, Y.; and Schapire, R. E. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55(1): 119–139.
- Gkoumas, D.; Uprety, S.; and Song, D. 2018. Investigating non-classical correlations between decision fused multimodal documents. In *International Symposium on Quantum Interaction*, 163–176. Springer.
- Glodek, M.; Tschechne, S.; Layher, G.; Schels, M.; Brosch, T.; Scherer, S.; Kächele, M.; Schmidt, M.; Neumann, H.; Palm, G.; and Schwenker, F. 2011a. Multiple Classifier Systems for the Classification of Audio-Visual Emotional States. In *Affective Computing and Intelligent Interaction*, 359–368. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Glodek, M.; Tschechne, S.; Layher, G.; Schels, M.; Brosch, T.; Scherer, S.; Kächele, M.; Schmidt, M.; Neumann, H.; Palm, G.; et al. 2011b. Multiple classifier systems for the classification of audio-visual emotional states. In *International Conference on Affective Computing and Intelligent Interaction*, 359–368. Springer.
- Hughes, R. I. 1989. *The structure and interpretation of quantum mechanics*. Harvard university press.
- Kolmogorov, A. N. 1950. Foundations of the theory of probability, Chelsea Pub. New York (German original 1933).
- Li, Q.; Gkoumas, D.; Lioma, C.; and Melucci, M. 2020. Quantum-inspired multimodal fusion for video sentiment analysis. *Information Fusion*.
- Li, Q.; Wang, B.; and Melucci, M. 2019. CNM: An Interpretable Complex-valued Network for Matching. 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)*, 4139–4148. Minneapolis, Minnesota: Association for Computational Linguistics. doi:10.18653/v1/N19-1420.
- Liang, P. P.; Liu, Z.; Bagher Zadeh, A.; and Morency, L.-P. 2018. Multimodal Language Analysis with Recurrent Multi-stage Fusion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 150–161. Brussels, Belgium: Association for Computational Linguistics. doi:10.18653/v1/D18-1014.
- Liang, Y.-C.; Spekkens, R. W.; and Wiseman, H. M. 2011. Specker’s parable of the overprotective seer: A road to contextuality, nonlocality and complementarity. *Physics Reports* 506(1-2): 1–39.
- Liu, Z.; Shen, Y.; Lakshminarasimhan, V. B.; Liang, P. P.; Bagher Zadeh, A.; and Morency, L.-P. 2018. Efficient Low-rank Multimodal Fusion With Modality-Specific Factors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2247–2256. Melbourne, Australia: Association for Computational Linguistics. doi:10.18653/v1/P18-1209.
- Melucci, M. 2008. A basis for information retrieval in context. *ACM Transactions on Information Systems (TOIS)* 26(3): 1–41.
- Melucci, M. 2015. *Introduction to information retrieval and quantum mechanics*. Springer.
- Morgenstern, O. 1949. The theory of games. *Scientific American* 180(5): 22–25.
- Morvant, E.; Habrard, A.; and Ayache, S. 2014. Majority Vote of Diverse Classifiers for Late Fusion. In *Structural, Syntactic, and Statistical Pattern Recognition - Joint IAPR International Workshop, S+SSPR 2014, Joensuu, Finland, August 20-22, 2014. Proceedings*, volume 8621 of *Lecture Notes in Computer Science*, 153–162. Springer.
- Nielsen, M. A.; and Chuang, I. L. 2011. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. USA: Cambridge University Press, 10th edition. ISBN 1107002176.

- Nojavanasghari, B.; Gopinath, D.; Koushik, J.; Baltrušaitis, T.; and Morency, L.-P. 2016. Deep multimodal fusion for persuasiveness prediction. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, 284–288.
- Pham, H.; Liang, P. P.; Manzini, T.; Morency, L.; and Póczos, B. 2019. Found in Translation: Learning Robust Joint Representations by Cyclic Translations between Modalities. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, 6892–6899. AAAI Press.
- Ponti Jr, M. P. 2011. Combining classifiers: from the creation of ensembles to the decision fusion. In *2011 24th SIBGRAPI Conference on Graphics, Patterns, and Images Tutorials*, 1–10. IEEE.
- Sordoni, A.; He, J.; and Nie, J.-Y. 2013. Modeling Latent Topic Interactions Using Quantum Interference for Information Retrieval. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13*, 1197–1200. New York, NY, USA: Association for Computing Machinery. ISBN 9781450322638.
- Sordoni, A.; Nie, J.-Y.; and Bengio, Y. 2013. Modeling Term Dependencies with Quantum Language Models for IR. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '13*, 653–662. New York, NY, USA: Association for Computing Machinery. ISBN 9781450320344.
- 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. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6558–6569. Florence, Italy: Association for Computational Linguistics.
- Tversky, A.; and Kahneman, D. 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review* 90(4): 293.
- Uola, R.; Luoma, K.; Moroder, T.; and Heinosaari, T. 2016. Adaptive strategy for joint measurements. *Physical Review A* 94(2): 022109.
- Uprety, S.; Tiwari, P.; Dehdashti, S.; Fell, L.; Song, D.; Bruza, P.; and Melucci, M. 2020. Quantum-like structure in multidimensional relevance judgements. *Advances in Information Retrieval* 12035: 728.
- Van Rijsbergen, C. J. 2004. *The geometry of information retrieval*. Cambridge University Press.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, 5998–6008.
- Wang, B.; Li, Q.; Melucci, M.; and Song, D. 2019a. Semantic Hilbert Space for Text Representation Learning. In *The World Wide Web Conference, WWW '19*, 3293–3299. New York, NY, USA: Association for Computing Machinery. ISBN 9781450366748. doi:10.1145/3308558.3313516.
- Wang, J.; Song, D.; and Kaliciak, L. 2010. Tensor Product of Correlated Textual and Visual Features: A Quantum Theory Inspired Image Retrieval Framework. In *2010 AAAI Fall Symposium Series*.
- Wang, Y.; Shen, Y.; Liu, Z.; Liang, P. P.; Zadeh, A.; and Morency, L. 2019b. Words Can Shift: Dynamically Adjusting Word Representations Using Nonverbal Behaviors. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 7216–7223. AAAI Press.
- Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; and Morency, L.-P. 2017. Tensor Fusion Network for Multimodal Sentiment Analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1103–1114. Copenhagen, Denmark: Association for Computational Linguistics. doi:10.18653/v1/D17-1115.
- Zadeh, A.; Liang, P. P.; Mazumder, N.; Poria, S.; Cambria, E.; and Morency, L. 2018. Memory Fusion Network for Multi-view Sequential Learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 5634–5641. AAAI Press.
- 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*.
- Zhang, P.; Niu, J.; Su, Z.; Wang, B.; Ma, L.; and Song, D. 2018a. End-to-End Quantum-like Language Models with Application to Question Answering. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, 5666–5673. AAAI Press.
- Zhang, Y.; Li, Q.; Song, D.; Zhang, P.; and Wang, P. 2019. Quantum-Inspired Interactive Networks for Conversational Sentiment Analysis. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 5436–5442. International Joint Conferences on Artificial Intelligence Organization.
- Zhang, Y.; Song, D.; Zhang, P.; Wang, P.; Li, J.; Li, X.; and Wang, B. 2018b. A quantum-inspired multimodal sentiment analysis framework. *Theoretical Computer Science* 752: 21 – 40. ISSN 0304-3975. Quantum structures in computer science: language, semantics, retrieval.