QeMMA: Quantum-Enhanced Multi-Modal Sentiment Analysis

Arpan Phukan[1] and **Asif Ekbal**[2]
Department of Computer Science and Engineering
Indian Institute of Technology Patna
Bihar
India, 801106
[1]arpanphukan@gmail.com, [2]asif@iitp.ac.in

Abstract

Multi-modal data analysis presents formidable challenges, as developing effective methods to capture correlations among different modalities remains an ongoing pursuit. In this study, we address multi-modal sentiment analysis through a novel quantum perspective. We propose that quantum principles, such as superposition, entanglement, and interference, offer a more comprehensive framework for capturing not only the cross-modal interactions between text, acoustics, and visuals but also the intricate relations within each modality. To empirically evaluate our approach, we employ the CMU-MOSEI dataset as our testbed and utilize Qiskit by IBM to run our experiments on a quantum computer. Our proposed Quantum-Enhanced Multi-Modal Analysis Framework **(QeMMA)** showcases its significant potential by surpassing the baseline by 3.52% and 10.14% in terms of accuracy and F1 score, respectively, highlighting the promise of quantum-inspired methodologies.

1 Introduction

Sentiments play a vital in our daily lives, serving as crucial elements in communication, decision-making, situational awareness, and learning. The commercial value of user-generated service and product reviews is substantial, as they aid fellow users in their decision-making processes, such as when purchasing a new product. Moreover, these reviews prove invaluable to businesses, offering insights into product monitoring, enhancing customer relationships, crafting more effective marketing strategies, and refining services.

Traditional systems typically rely on a single modality, text (Akhtar et al., 2016, 2017; Zadeh et al., 2018a), or other source (Phukan and Gupta, 2022b,a, 2023) to gauge a person's sentiment or emotional state. Conversely, multi-modal Sentiment Analysis seeks to empower machines with a deeper understanding of sentiment by leveraging information from multiple modalities, including text, visuals, and acoustics (Wu et al., 2022; Huddar et al., 2021). Video content, in particular, serves as a rich source for extracting multi-modal information. In addition to visual frames, videos provide acoustic and textual imitation of verbal language. Furthermore, a single video may contain multiple utterances from a single speaker, each potentially expressing different sentiments. Thus, the sentiment polarity of one statement often hinges on contextual utterances, creating intricate interdependencies. Prior studies in multi-modal sentiment analysis have explored various approaches, such as applying attention mechanisms to contextual utterances for classification (Poria et al., 2017), computing correlations among modalities in target and contextual utterances (Ghosal et al., 2018), and dynamically adjusting word representations based on nonverbal cues via LSTMs (Wang et al., 2019). However, these methods are computationally intensive, imposing significant overhead for only marginal gains in accuracy.

In this paper, we address this challenge by introducing Quantum Machine Learning. Our novel approach combines a Variational Quantum Circuit (VQC) (Qi et al., 2021) with a recurrent neural network-based multi-modal multi-utterance attention architecture for sentiment analysis. We propose the following:

Quantum Superposition (QS): Expressing Intra-Modal Affinities: *QS, where a quantum state simultaneously occupies numerous mutually incompatible ground states with a probability distribution until measured*, offers an apt analogy for multi-modal tasks. We hypothesize that an utterance can be viewed as a QS of various text-to-text, audio-to-audio, or visual-to-visual instances, yielding complex probability amplitudes suitable for quantum probability theory.

Quantum Interference and Entanglement:

Multi-Modal Feature Fusion: *Quantum Interference (QI), embodying the interference of two propagation paths (e.g., acoustics and textual channels), influences the probability distribution of a system (e.g., the sentiment polarity)*. In a similar fashion, *Quantum Entanglement (QE) captures the correlation of particles (modalities) where the state of a particle is dependent on the other, irrespective of their distance*. Introducing QI and QE terms defines a non-linear fusion mechanism, mirroring the characteristic where each modality contributes simultaneously to the multi-modal system, resulting in inter-modal dependence.

Quantum Measurement (QM): Reflecting Correlations Across Modalities: The three modalities (text, visual, and acoustic) are inseparable from each other. For instance, a seemingly positive sentence like *"I love to take a cold shower at 3 am in the morning during winters!!"* gains negative connotations when supplemented with video evidence (unpleasant gestures) and acoustic cues (intensity, pitch). Thus, text, visual, and acoustic features together facilitate the correct sentiment identification. This correlation can be effectively modeled by *QM, where measuring a qubit's state to reflect over a classical bit perturbs the states of the other qubits in the quantum system*. These qubits (modalities) are intimately linked as each modality plays a part in correctly interpreting the sentiment label.

The primary motivation of our study is to harness quantum characteristics, including superposition, interference, and entanglement, to replicate both inter and intra-modal relationships. Our contributions in this paper are as follows: **(1)** We introduce a novel VQC designed to leverage quantum properties for capturing multi-modal interactions among text, acoustics, and visuals. We propose that QS naturally encompasses intra-modal associations, while entanglement inherently models the influence of the three modalities on each other. **(2)** Building on the work of Ghosal et al. (2018), we introduce quantum elements into their proposed architecture to facilitate a fair comparison and illustrate the potential of quantum-enhanced approaches in multi-modal sentiment analysis.

2 Related Work

Quantum Natural Language Processing (QNLP): In QNLP, two primary perspectives have emerged. First, there is a theoretical application of quantum probability to address classification (Zhang et al., 2020) and generation tasks (Zhang et al., 2018a) in classical computing systems. Second, there is a practical implementation angle, wherein natural language tasks are executed on quantum simulators (Li et al., 2023) and actual quantum computers (Qi et al., 2021). Of note, the Categorical Distributional Compositional (DisCoCat) model for natural language, leveraging quantum instantiable algorithms, has gained prominence due to the growing availability of quantum resources. In a recent study conducted by Meichanetzidis et al. (2023), the DisCoCat model was scrutinized for sentence classification. They instantiated sentences and word meanings as parameterized quantum circuits, with quantum entanglement being used to manifest the grammatical structure. Additionally, Lorenz et al. (2021) delved into the mapping of sentence representations into quantum circuits, exploring syntax-independent models like the "bag-of-words" paradigm in the quantum realm. Researchers have also endeavored to theoretically establish quantum probability as an effective means of fusing different modalities (Zhang et al., 2018b).

Sentiment Analysis: Multi-modal sentiment prediction has gained popularity in recent years, with researchers experimenting with various methodologies for merging distinct modalities. Hybrid fusion methods, which concurrently scrutinize feature-level and decision-level attributes, have garnered attention (Poria et al., 2015, 2016). Tensor fusion approaches, exemplified by T2FN (Liang et al., 2019), and MTFN-CMM (Yan et al., 2022) involve the construction of a 3-fold Cartesian product through the utilization of modality embeddings. These methods explicitly model unimodal, bimodal, and trimodal interactions using a tensor fusion layer. For word-level fusion, LSTM-based gated modality mixing networks, such as RMFN (Liang et al., 2018), and RAVEN (Wang et al., 2019) rely on bidirectional recurrent neural networks. Meanwhile, multi-head attention-based networks like BIMHA (Wu et al., 2022) and MMHA (Xi et al., 2020) extract contextual features using a bidirectional recurrent neural network-based model.

3 Proposed Methodology

To obtain quantum representations of text, visual, and acoustic inputs, their dimension needs to be reduced due to constraints in qubit availability on

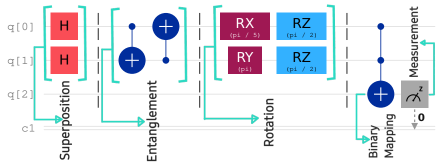

Figure 1: Our Proposed VQC.

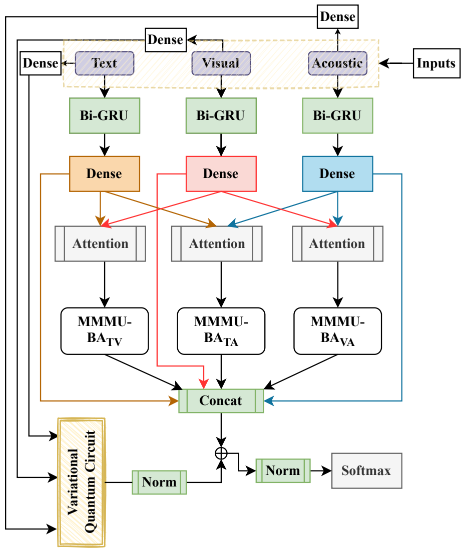

Figure 2: QeMMA. Here, $\bigoplus$ is an addition operation.

noisy intermediate-scale quantum computers (Qi et al., 2021). Separate dense layers are applied to each modality, reducing the dimension to 1. Subsequently, these modalities are introduced as quantum trainable parameters φ_0, φ_1, and φ_2 influencing qubit rotations in our VQC, detailed in Section 3.1. The VQC output is then integrated into the baseline architecture.

To ensure a transparent and equitable comparison, we adopt the MMMU-BA model architecture proposed by Ghosal et al. (2018). Their framework is designed to process multi-modal information, encompassing text, visual, and acoustic data across a sequence of utterances. This information is channeled into three distinct Bi-GRU layers. Subsequently, a dense operation is applied uniformly across the time steps, with dedicated layers for each modality. Following this, they apply multi-modal attention to the outputs of these dense layers. Their approach involves implementing a bi-modal attention framework, where an attention mechanism is employed to weigh the contributions of pairwise modalities, namely text-acoustic, visual-text, and acoustic-visual.

The results from these pairwise attention mechanisms, combined with modality representations, are concatenated and added to the VQC output. This combined representation is then passed through a softmax layer for classification. We illustrate the model architecture in Figure 2.

The execution of our QeMMA model is carried out using the **qasm_simulator** from the **Aer backend** and the **ibmq_lagos quantum computer** via the **qiskit** library.

3.1 QeMMA: Quantum-Enhanced Multi-Modal Analysis Framework

Our methodology for incorporating the quantum algorithm into the baseline is as follows: **(a)** To ensure a fair comparison, we meticulously adhere to the pre-processing and classical model conventions established by the baseline framework. **(b)** Independent dense layers are employed to adjust the dimensions of the three classical input modalities, viz. text, visual and acoustic. This transformation ensures compatibility with the number of trainable parameters in the VQC. **(c)** The VQC introduced to the system facilitates superposition, entanglement, rotation, and measurement of the qubits. This rotation is governed by the values of specific parameters, yielding an output state represented as $\left|\Psi_{out}(x_i, \varphi_0, \varphi_1, \varphi_2)\right\rangle = \Upsilon(\varphi_0, \varphi_1, \varphi_2)\left|\Psi_{in}(x_i)\right\rangle$, where $\Upsilon(\varphi_0, \varphi_1, \varphi_2)$ corresponds to the VQC, $\Psi_{out}(x_i, \varphi_0, \varphi_1, \varphi_2)$ refers to the output state of the quantum system, and $\Psi_{in}(x_i)$ is the input to the quantum system. **(d)** We employ a Toffoli gate to account for our binary classification task (Refer to Section 3.2). We then perform the QM operation on the circuit to determine the expectation values. This operation serves to collapse the quantum state into classical values. **(e)** Our iterative approach involves tuning φ_p where $p \in \{0, 1, 2\}$ to minimize the cost function. This optimization process is facilitated through the utilization of the parameter shift rule (Mitarai et al., 2018).

3.2 Variational Quantum Circuit (VQC)

In Figure 1, we illustrate our proposed VQC, which constitutes the core of our quantum-enhanced methodology. Within this framework, we employ a set of qubits labeled as q_i, with $i \in \{0, 1, 2\}$. The Hadamard gates, illustrated as H blocks, play a pivotal role in inducing superposition for qubits q_0 and q_1. Subsequently, to create entanglement within the system, we introduce two controlled-NOT gates. After establishing superposition and entanglement among the qubits, we execute rotation operations along the X, Y, and Z axes. The R_x block signifies the Pauli rotation X gate, enabling qubit rotation as described by the equation $R_x(\varphi_0) = \begin{bmatrix} \cos\frac{\varphi_0}{2} & -i\sin\frac{\varphi_0}{2} \\ -i\sin\frac{\varphi_0}{2} & \cos\frac{\varphi_0}{2} \end{bmatrix}$. Similarly, the R_y block corresponds to Pauli rotation Y gate, in accordance with the equation $R_y(\varphi_1) = \begin{bmatrix} \cos\frac{\varphi_1}{2} & -\sin\frac{\varphi_1}{2} \\ \sin\frac{\varphi_1}{2} & \cos\frac{\varphi_1}{2} \end{bmatrix}$. Lastly, the R_z block designates the integration of Pauli rotation Z gates, facilitating qubit rotation by $R_z(\varphi_2) = \begin{bmatrix} e^{-i\varphi_2} & 0 \\ 0 & e^{\varphi_2} \end{bmatrix}$. Subsequently, we introduce a Toffoli gate to alter the polarity of qubit q_2 based on the states of qubits q_0 and q_1. Thereafter, we measure the state of qubit q_2 and map it to a classical bit. This strategy of implementing a Toffoli gate and mapping q_2 to a classical bit is adopted because the task primarily entails binary classification, with the QM of a single qubit yielding possible states of 0 and 1. Moreover, this approach helps mitigate the impact of quantum noise to some extent. By reducing the number of potential qubit states, we effectively minimize the potential for erroneous state readings due to quantum noise. This reduction ensures that the state probability of the entire system (calculated as $\sum_0^{i=2^N} S_i \times P_i$, where S_i is the possible state, P_i is the probability of the quantum system collapsing into state S_i and N = Total number of qubits) is not adversely affected by noisy states and their probabilities, thereby ensuring the robustness of our quantum classification approach.

4 Datasets, Experiments and Results

4.1 Datasets

Our study utilizes the CMU-MOSEI dataset Zadeh et al. (2018b), comprising 3,229 videos and 22,676 utterances from over 1,000 YouTube speakers. Training, validation, and test sets contain 16,216, 1,835, and 4,625 utterances, respectively. Sentiment labels originally range from -3 to +3 but are transformed into a two-class system (positive and negative) for consistency with the baseline.

4.2 Experimental Setup

We have adopted the experimental configuration proposed by Ghosal et al. (2018) to ensure a consistent benchmark. To enhance model generalization and combat overfitting, they introduce dropout with a rate of 0.3 in both the dense layers and the BiGRU layers. The dense layers employed the ReLU activation function, while the final classification layer used a softmax activation. In the training process, they opt for the Adam optimizer in conjunction with a cross-entropy loss function. We set the epochs to 50, and as for our VQC, we specify a shift parameter of 0.5 and configure the quantum system to perform 1024 shots.

4.3 Results and Analysis

Table 1 provides a discerning overview of our results, demonstrating that QeMMA is comparable to a recent state-of-the-art system despite not surpassing it. Notably, our findings reveal a substantial 3.52% accuracy improvement and a remarkable 10.14% increase in the F1 score compared to the baseline MMMU-BA solely through incorporating the multi-modal quantum representation of the classical inputs. This outcome substantiates our initial hypothesis, indicating that quantum properties effectively emulate the intricate and hidden intra and inter-modal relationships. Note that this performance improvement is achieved with minimal overhead. QeMMA introduces a mere 809 additional parameters, with MMMU-BA housing 34,32,302 parameters, while the former has 34,33,111. This modest parameter addition, which yields significant performance gains, underscores the efficiency and reliability of our quantum-augmented approach.

We've listed the error cases of QeMMA and their potential reasons in Table 3. Further information regarding the error analysis can be found in the Appendix A.2.

Model	Acc_2	F1
Ghosal et al.'s (2018) MMMU-BA*	77.91	71.37
Koromilas and Giannakopoulos's (2021) CAE-LR	78.00	76.30
Paraskevopoulos et al.'s (2022) MMLATCH	81.20	81.41
QeMMA (OURS)	**81.43**	**81.54**

Table 1: Main Results. *MMMU-BA is the Baseline.

5 Conclusion

We introduced a novel VQC to the MMMU-BA baseline, leveraging quantum properties (superposition, entanglement, and interference). This enhanced performance by just 309 parameters, highlighting the efficiency of our quantum approach.

Furthermore, it is worth noting that the choice of the specific quantum circuit itself is a hyperparameter, suggesting potential optimization for specific tasks. In the future, we aim to refine our VQC and architecture to enhance multi-modal sentiment analysis with quantum capabilities.

References

Md Shad Akhtar, Deepak Gupta, Asif Ekbal, and Pushpak Bhattacharyya. 2017. Feature selection and ensemble construction: A two-step method for aspect based sentiment analysis. *Knowledge-Based Systems*, 125:116–135.

Md Shad Akhtar, Ayush Kumar, Asif Ekbal, and Pushpak Bhattacharyya. 2016. A hybrid deep learning architecture for sentiment analysis. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 482–493.

Deepanway Ghosal, Md Shad Akhtar, Dushyant Chauhan, Soujanya Poria, Asif Ekbal, and Pushpak Bhattacharyya. 2018. Contextual inter-modal attention for multi-modal sentiment analysis. In *proceedings of the 2018 conference on empirical methods in natural language processing*, pages 3454–3466.

Mahesh G Huddar, Sanjeev S Sannakki, and Vijay S Rajpurohit. 2021. Attention-based multimodal contextual fusion for sentiment and emotion classification using bidirectional lstm. *Multimedia Tools and Applications*, 80:13059–13076.

Panagiotis Koromilas and Theodoros Giannakopoulos. 2021. Unsupervised multimodal language representations using convolutional autoencoders. *arXiv preprint arXiv:2110.03007*.

Qiuchi Li, Benyou Wang, Yudong Zhu, Christina Lioma, and Qun Liu. 2023. Adapting pre-trained language models for quantum natural language processing.

Paul Pu Liang, Zhun Liu, Yao-Hung Hubert Tsai, Qibin Zhao, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2019. Learning representations from imperfect time series data via tensor rank regularization. *arXiv preprint arXiv:1907.01011*.

Paul Pu Liang, Ziyin Liu, Amir Zadeh, and Louis-Philippe Morency. 2018. Multimodal language analysis with recurrent multistage fusion. *arXiv preprint arXiv:1808.03920*.

R Lorenz, A Pearson, K Meichanetzidis, D Kartsaklis, and B Coecke. 2021. Qnlp in practice: Running compositional models of meaning on a quantum computer (2021). *URL https://arxiv. org/abs/2102.12846*.

Konstantinos Meichanetzidis, Alexis Toumi, Giovanni de Felice, and Bob Coecke. 2023. Grammar-aware sentence classification on quantum computers. *Quantum Machine Intelligence*, 5(1):10.

Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. 2018. Quantum circuit learning. *Physical Review A*, 98(3):032309.

Georgios Paraskevopoulos, Efthymios Georgiou, and Alexandras Potamianos. 2022. Mmlatch: Bottom-up top-down fusion for multimodal sentiment analysis. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4573–4577. IEEE.

Arpan Phukan and Deepak Gupta. 2022a. Eeg based emotion classification using xception architecture. In *Mobile Radio Communications and 5G Networks: Proceedings of Second MRCN 2021*, pages 95–108. Springer.

Arpan Phukan and Deepak Gupta. 2022b. Emotion classification using xception and support vector machine. In *Advanced Machine Intelligence and Signal Processing*, pages 721–734. Springer.

Arpan Phukan and Deepak Gupta. 2023. Deep feature extraction from eeg signals using xception model for emotion classification. *Multimedia Tools and Applications*, pages 1–19.

Soujanya Poria, Erik Cambria, and Alexander Gelbukh. 2015. Deep convolutional neural network textual features and multiple kernel learning for utterance-level multimodal sentiment analysis. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 2539–2544.

Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Mazumder, Amir Zadeh, and Louis-Philippe Morency. 2017. Multi-level multiple attentions for contextual multimodal sentiment analysis. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 1033–1038. IEEE.

Soujanya Poria, Erik Cambria, Newton Howard, Guang-Bin Huang, and Amir Hussain. 2016. Fusing audio, visual and textual clues for sentiment analysis from multimodal content. *Neurocomputing*, 174:50–59.

Jun Qi, Chao-Han Huck Yang, and Pin-Yu Chen. 2021. Qtn-vqc: An end-to-end learning framework for quantum neural networks.

Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. 2019. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7216–7223.

Ting Wu, Junjie Peng, Wenqiang Zhang, Huiran Zhang, Shuhua Tan, Fen Yi, Chuanshuai Ma, and Yansong Huang. 2022. Video sentiment analysis with bimodal information-augmented multi-head attention. *Knowledge-Based Systems*, 235:107676.

Chen Xi, Guanming Lu, and Jingjie Yan. 2020. Multimodal sentiment analysis based on multi-head attention mechanism. In *Proceedings of the 4th international conference on machine learning and soft computing*, pages 34–39.

Xueming Yan, Haiwei Xue, Shengyi Jiang, and Ziang Liu. 2022. Multimodal sentiment analysis using multi-tensor fusion network with cross-modal modeling. *Applied Artificial Intelligence*, 36(1):2000688.

Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, and Louis-Philippe Morency. 2018a.

Multi-attention recurrent network for human communication comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018b. 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)*, pages 2236–2246.

Peng Zhang, Jiabin Niu, Zhan Su, Benyou Wang, Liqun Ma, and Dawei Song. 2018a. End-to-end quantum-like language models with application to question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Yazhou Zhang, Dawei Song, Xiang Li, Peng Zhang, Panpan Wang, Lu Rong, Guangliang Yu, and Bo Wang. 2020. A quantum-like multimodal network framework for modeling interaction dynamics in multiparty conversational sentiment analysis. *Information Fusion*, 62:14–31.

Yazhou Zhang, Dawei Song, Peng Zhang, Panpan Wang, Jingfei Li, Xiang Li, and Benyou Wang. 2018b. A quantum-inspired multimodal sentiment analysis framework. *Theoretical Computer Science*, 752:21–40.

A Appendix

A.1 Dataset

We present the statistics of the dataset in Table 2.

Statistics	CMU-MOSEI		
	Tr	Dv	Ts
Videos	2250	300	679
Utterance	16216	1835	4625
Utterance/Video - Min	1	1	1
Utterance/Video - Max	98	37	52
Utterance/Video - Avg	7.207	6.116	6.821
Positive	11498	1332	3281
Negative	4718	503	1344
Words/Utter. - Min	1	1	1
Words/Utter. - Max	515	224	549
Words/Utter. - Avg	18.227	18.498	18.658
Utter-Len/Video - Min	0.089s	0.22s	0.15s
Utter-Len/Video - Max	208.27s	90.42s	188.22s
Utter-Len/Video - Avg	6.896s	6.960s	7.158s
Speakers	1000		

Table 2: Dataset Statistics

A.2 Error Analysis

Table 3 highlights an interesting observation where our model struggles to classify the sentiment of the statement, "IS IT GONNA BE GOOD." This difficulty may stem from the statement's inclination towards neutrality. The speaker expresses uncertainty about the movie's quality, driven by a mixture of anticipation (positive sentiment) and hesitation (negative sentiment) due to the presence of robots. Although the context offers sufficient clues to label the utterance as negative, QeMMA, in this instance, fails to do so.

Transcript	Sentiment	Predicted
RELEASE THE CRACKEN ON THIS MOVIE DONT GO SEE IT GO ... SAY ABOUT THAT	Negative	Negative
SO DISAPPOINTMENT DEFINATELY	Negative	Negative
I DID NOT WANNA GIVE THIS MOVIE A BAD REVIEW	Negative	Negative
I IM A FAN OF THE ORIGINAL YOU KNOW LOVED IT	Positive	Positive
UM I WAS REALLY LOOKING FORWARD TO IT	Negative	Negative
AND UM BUT I GOT TO I GOT TO DO MY DUTY AND LET YOU GUYS KNOW ITS A A MAJOR SUCK FEST	Negative	Negative
HEARD GOOD THINGS ABOUT THE SCREENINGS ON THAT	Positive	Positive
AND UM LOOKING REALLY FORWARD TO SEEING THAT MOVIE	Positive	Positive
I WAS A LITTLE HASITANT AT FIRSTYOU KNOW AS YOU ARE WITH WITH REBOOTS YOU KNOW	Negative	Negative
IS IT GONNA BE GOOD	**Negative**	**Positive**
BUT AS THE TRAILERS CAME OUT I I REALLY LIKED WHAT I WAS SEEING	Positive	Positive
I THINK THAT ... ME EXCITED ABOUT SEEING THIS MOVIE	Positive	Positive
AND AND BEING A FAN OF THE ORIGINAL I WAS I WAS LOOKING FORWARD TO SEEING IT	Negative	Negative
BUT TO ME SOMEONE NEEDS TO RELEASE THE CRACKON ON THIS MOVIE AND JUST TOTALLY DESTROY IT	Negative	Negative

Table 3: Error Analysis