

# From None to Severe: Predicting Severity in Movie Scripts

Yigeng Zhang<sup>†</sup>, Mahsa Shafaei<sup>†</sup>, Fabio A. González<sup>‡</sup> and Thamar Solorio<sup>†</sup>

<sup>†</sup>University of Houston

<sup>‡</sup>Universidad Nacional de Colombia

<sup>†</sup>{yzhang168, mshafaei, tsolorio}@uh.edu

<sup>‡</sup>fagonzalezo@unal.edu.co

## Abstract

In this paper, we introduce the task of predicting severity of age-restricted aspects of movie content based solely on the dialogue script. We first investigate categorizing the ordinal severity of movies on 5 aspects: *Sex*, *Violence*, *Profanity*, *Substance consumption*, and *Frightening scenes*. The problem is handled using a Siamese network-based multitask framework which concurrently improves the interpretability of the predictions. The experimental results show that our method outperforms the previous state-of-the-art model and provides useful information to interpret model predictions. The proposed dataset and source code are publicly available at [our GitHub repository](#)<sup>1</sup>.

## 1 Introduction

Estimating the severity level of objectionable content for movies can provide convenience for users to judge whether a movie is suitable for watching. For example, parents may want to make sure there are no violent scenes in a movie when they plan to watch it with their kids, because exposure to violent scenes may increase youth aggressive behavior and decrease their empathy (Anderson et al., 2017). However, existing rating systems (e.g., MPAA) only provide simple age restrictions and do not include the suitability level on a specific aspect of the content. Furthermore, a system that can automatically track the severity level of objectionable content helps the creative professionals evaluate the age suitability of their work. They may get assisted by this function and adjust the product creation or marketing strategy based on the corresponding target audiences. This system can be easily applied to any dialogue-intensive compositions like novel and screenplay writing. Content evaluation and intervention by the writers can happen at any stage of production to assess age-restricted contents.

In this work, we propose to solve the problem of predicting the severity of age-restricted content solely using the dialogue script data. Text is much more lightweight than visual data (such as images and videos), so the processing procedure can be more efficient and scalable considering the increasing fidelity of multimedia content. We initiate our exploration on movies from five aspects of contents: *Sex & Nudity*, *Violence & Gore*, *Profanity*, *Alcohol*, *Drugs & Smoking*, and *Frightening & Intense Scenes* as used in *IMDB*<sup>2</sup> Parent Guide.

There are a small number of previous works that studied modeling age-restricted content. (Shafaei et al., 2020) initiated the research of predicting MPAA ratings of the movies leveraging movie script and metadata. (Martinez et al., 2019) focused on violence detection using movie scripts while (Martinez et al., 2020) expanded the scope to violence, substance abuse, and sex. Both works intended to predict the severity of age-restricted content into three manually defined levels: low, mid, and high. In this work, We introduce two more aspects of interest: *Frightening* and *Profanity*. Instead of manually downgrading severity levels into three categories, we explore with a more challenging setting: rating on 4 originally defined fine-grained severity levels from collective rating by customers: *None*, *Mild*, *Moderate*, and *Severe*.

The major contributions of our research can be summarized as follows: (1) This work is the first attempt to solve the age-restricted content severity predicting problem from 5 aspects. We studied multiple baselines and presented a competitive method to inspire future exploration. (2) To our best knowledge, the dataset we developed is the first publicly available dataset for this task. The size is roughly five times larger than the restricted datasets from the previous works. (3) We proposed an effective multitask ranking-classification framework to solve

<sup>1</sup><https://github.com/RiTUAL-UH/Predicting-Severity-in-Movie-Scripts>.

<sup>2</sup><https://www.imdb.com/>, one of the most visited online databases of film, TV and celebrity content.

this problem. Our method dealt with long movie scripts successfully and achieved state-of-the-art results in all five aspects of age-restricted contents with rich interpretability.

## 2 Methodology

We investigate predicting the severity of 5 objectionable aspects of movies. This problem is formulated as a multi-class classification task. The average length of the dialogue scripts is around 10,000 words, which drastically exceeds the limit of current popular Transformer-based models. To leverage the strong semantic representation capability of the Transformers, we propose to represent each utterance as the basic unit, and further encode the context with the recurrent modules. Finally, we use a fully connected layer on top of the encoded representations to produce the classification predictions.

For this model, we first leverage SentenceTransformers (Reimers and Gurevych, 2019) to encode each dialogue utterance. Then, a Bi-directional LSTM encoder is deployed to model the sequential interrelations of the utterance flow. We finally apply a max-pooling operation on all time steps of the hidden states of the recurrent module to get the document representation for classification following the practice in (Howard and Ruder, 2018). We also study another strong word-level deep learning model, TextRCNN (Lai et al., 2015), to probe the significance of lexical signals.

### 2.1 The Multitask Ranking-Classification Framework

The severity of a particular age-restricted content is a relative concept. People assign severity ratings to movies based on their own experiences and personal beliefs. Meanwhile, the severity levels are ordinal variables instead of independent categorical classes. Therefore, customers can gain a vivid understanding of the severity levels of an unfamiliar movie when comparing it to some examples (e.g., previous watched ones).

The general model development algorithm is described as shown in Algorithm 1. We assume that learning to compare movies on their severity is a proxy to understand how the model differentiates severity. So we propose a pairwise ranking objective (Hüllermeier et al., 2008) as an auxiliary task to probe into the model behavior for a more interpretable prediction. Other than existing multitask

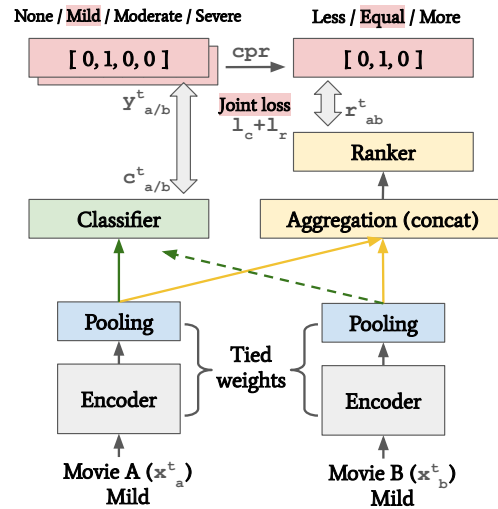


Figure 1: Multitask pairwise ranking-classification network.

practices for text classification (Liu et al., 2016; Zhang et al., 2017), this framework is based on a Siamese network with tied-weights for both instance classification and comparison, and it can be adopted by any backbone encoder. Here we apply it to the Bi-LSTM + Transformers model (RNN-Trans) and TextRCNN model. The classification and the ranking objectives learn over two individual cross-entropy losses. Then the model is optimized on a joint loss of classification and ranking. The model structure is illustrated in Figure 1.

---

#### Algorithm 1: Ranking-Classification

---

**Input:** training instance set  $\mathbb{X}_t$  with severity label set  $\mathbb{Y}_t$ , ranking-classification model  $f$ , comparison operation  $cpr$ , classification/ranking loss  $\mathcal{L}_c/\mathcal{L}_r$ .

**Output:** multitask ranking-classification model  $\hat{f}$

```

1 Function RANK-CLS ( $f, \mathbb{X}_t, \mathbb{Y}_t$ ):
2   initialization;
3   while not stopping criteria do
4     randomly pick  $x_i^t, x_j^t \in \mathbb{X}_t$  with
       corresponding  $y_i^t, y_j^t \in \mathbb{Y}_t$ ;
5      $c_{i,j}, r_{ij} \leftarrow f(x_i^t, x_j^t)$ ;
6      $l_c \leftarrow \mathcal{L}_c(c_{i,j}, y_{i,j}^t)$ ;
7      $l_r \leftarrow \mathcal{L}_r(r_{ij}, cpr(y_i^t, y_j^t))$ ;
8      $\hat{f} \leftarrow \underset{f}{\operatorname{argmin}} (l_c + l_r)$ ;
9   end
10 return

```

---

Method		Sex	Violence	Profanity	Substance	Frightening	Avg
Baseline model	LR + Glove Avg.	33.87	46.35	48.06	29.27	41.38	39.79
	SVM + Glove Avg.	27.48	41.88	44.16	18.68	35.42	33.52
	TextCNN	38.19	46.61	63.95	37.16	44.82	46.14
	BERT	33.73	36.29	49.48	34.58	37.75	38.37
Backbone model	TextRCNN	43.13	52.51	66.49	41.79	49.74	50.73
	RNN-Trans	44.76	55.72	62.32	42.39	50.95	51.23
Proposed multi-task model	TextRCNN S-MT	43.01	52.59	<b>67.26</b>	<b>43.92</b>	50.36	51.43
	RNN-Trans S-MT	<b>45.68</b>	<b>55.90</b>	62.65	42.21	<b>51.55</b>	<b>51.60</b>

Table 1: Severity prediction macro F1 scores on test data.

The auxiliary ranking component will distinguish the severity difference out of the training pairs, with a learning objective of predicting lower/equal/higher severity between the two instances. By introducing this function, we can apply model  $f$  to compare the severity level of any pair of movies given an aspect.

### 3 Dataset

The dataset used in this work was developed based on the script data used in (Shafaei et al., 2019, 2020). We collected the up-to-date user ratings for age-restricted content from *IMDB.com* for more than 15,000 movies. The age-restricted aspects are adopted from the *Parents Guide* section of each movie, and there are five aspects: *Sex & Nudity*, *Violence & Gore*, *Profanity*, *Alcohol*, *Drugs & Smoking*, and *Frightening & Intense Scenes*. Each of the aspects has four severity levels for the users to rate on the corresponding movies from low to high, which are *None*, *Mild*, *Moderate*, and *Severe*. In this work, we pick the ratings on the website as the category label for each aspect.

After collecting the user ratings, we filter out movies with less than 5 votes to make sure the collected severity level is robust to use. At last, we have roughly 4,400 to 6,600 movies for each aspect.

The movie scripts have a median/average length of around 10,000 words. The vocabulary size of each aspect is roughly 330,000 to 450,000. Detailed dataset descriptions are attached in the appendix.

Comparing to the previous works with restricted data access (Martinez et al., 2019, 2020), our dataset is roughly five times larger, and the data source is free to access. We made the updated dataset publicly available with the same data partitions in this work for reproducibility purposes.

Method	Sex.	Vio.	Pro.	Sub.	Fri.
(Shafaei et al., 2020)	29.21	36.65	50.57	33.48	27.82
(Martinez et al., 2020)	40.91	53.02	60.51	35.60	48.81
TextRCNN S-MT	41.27	54.11	<b>69.51</b>	<b>43.56</b>	47.18
RNN-Trans S-MT	<b>44.66</b>	<b>55.29</b>	64.01	42.63	<b>51.03</b>

Table 2: Performance benchmarking with related approaches over the same 10-fold cross-validation.

### 4 Experimental Results

We evaluate the effectiveness of different methods on macro F1 score because the data is imbalanced. The dataset is first split and stratified using an 80/10/10 ratio for training, development, and test for each age-restricted aspect. The experimental results on the test set are reported in Table 1. Baseline models include average GloVe embedding (Pennington et al., 2014) with SVM/Logistic Regression, TextCNN (Kim, 2014), and BERT (Devlin et al., 2019). The proposed multitask model outperforms multiple baselines by a compelling margin in all aspects. Statistical significance test shows introducing the ranking subtask does no harm to model performance.

The proposed RNN-Trans multitask Siamese model (RNN-Trans S-MT) dominates on *Sex & Nudity*, *Violence & Gore*, and *Frightening & intensive scenes* while the proposed TextRCNN multitask Siamese model (TextRCNN S-MT) works best in *Profanity* and *Alcohol*, *Drugs & Smoking*. This outcome is intuitive because *Profanity* is the only aspect that has an overt pattern such as bad words and abusive language included in the dialogue script. Those utterances are neither missing nor latent to both audiences and the NLP model. So word-level models can have advantages over the utterance-level model in catching such signals. As for *Frightening & intensive scenes* and *Violence & Gore*, they are more challenging for the model to make inference because the data lacks visual and audio information. Dialogues might sometimes imply scenes such as a violent fight with cursing

Movie	Aspect	Gold label	Prediction	None	Mild	Moderate	Severe
Deadpool (2016)	Profanity	Severe	Severe ✓	●●●●●●	●●●●●●	●●●●●●	●●●●●●
Pride & Prejudice (2005)	Violence	None	None ✓	●●●●●●	●●●●●●	●●●●●●	●●●●●●
Django Unchained (2012)	Sex	Moderate	Mild ✗	●●●●●●	●●●●●●	●●●●●●	●●●●●●
A Clockwork Orange (1971)	Substance	Moderate	Mild ✗	●●●●●●	●●●●●●	●●●●●●	●●●●●●
The Greatest Showman (2017)	Frightening	Mild	Mild ✓	●●●●●●	●●●●●●	●●●●●●	●●●●●●

Table 3: Analysis on successful/unsuccessful test examples. Predictions with ✓ mean correct while with ✗ represents incorrect. For candidate comparators in each severity level, a black circle indicates the model believes the test sample has a higher severity level than the comparator. Similarly, a gray circle means the test sample has an equal severity level, and a light gray circle means lower severity. The comparison results come from best-performing models of each aspect.

words or a frightening shot with screaming, however, not all such scenes come with particular dialogues. For *Sex & Nudity* and *Alcohol, Drugs & Smoking*, it is even more difficult to infer from plain text without any visual evidence.

We also experiment with models from related previous works by (Shafaei et al., 2020; Martinez et al., 2020). The former leverages LSTM with attention and NRC emotion to model the textual signal to predict MPAA ratings for the movie with the script; the latter uses RNN on dialogue encoded by a pretrained MovieBERT model and sentiment embeddings to predict severity of different aspects simultaneously. For fair and reliable comparison, we conduct the benchmarking within the same setting of using the dialogue script as the only available information in the same 10-fold cross-validation configuration from (Martinez et al., 2020). Table 2 shows our proposed method outperformed previous works by a large margin.

To conclude, our proposed method works reliably better in *Frightening & intensive scenes* and *Violence & Gore* due to Transformer’s strong ability to capture latent implications behind utterances. While the word-level model, TextRCNN, performs better in detecting overt signals in *Profanity*. For *Sex & Nudity* and *Alcohol, Drugs & Smoking*, our method still achieved the state-of-the-art although modeling the severity for these aspects remains challenging.

## 5 Discussion and Analysis

We investigate several popular movies with at least 200,000 IMDB ratings from the test set of each aspect. We collect the top 5 movies with the largest number of severity ratings from each severity level as comparators, then do the pairwise ranking between each movie and all its comparators for test.

Comparison results are shown in Table 3.

**Successful examples:** For the movie *Deadpool (2016)*, the model gives a correct prediction on the *Profanity* aspect with a convincing signal from pairwise ranking. This movie is determined to have a higher severity level in profanity than any candidate comparators from *None/Mild/Moderate* level. For *Severe*-level comparators, the test instance is lower than one, equal to three, and higher than one of the comparators. The movie *Pride & Prejudice (2005)* comes with a correct prediction on *Violence*. Four comparators from *None* show they have equal severity while only one denotes higher. Three *Mild* comparators give equal results while two indicate the test movie is lower than them. For moderate and severe comparators, all of them evince the test movie has lower severity. We can therefore reason that *Pride & Prejudice (2005)* has a severity level of *None* in terms of violence content. There is a similar case as presented in *The Greatest Showman (2017)*. It tends to appear higher than none but roughly *Mild* level among the comparators. Then the model gives it a correct prediction to *Mild*. Therefore, by investigating the pairwise ranking results, the interpretation of this prediction is more vivid to make sense to users concerning the relative severity level.

**Unsuccessful examples:** For the movie *Django Unchained (2012)* on *Nudity* and the movie *A Clockwork Orange (1971)* on *Alcohol, Drugs & Smoking*, the ranking gives a very clear pattern: higher than none, equal to mild, and lower than moderate. However, the actual severity level is moderate instead of mild. We conservatively argue aspects such as *Nudity* and *Alcohol, Drugs & Smoking* can hardly be inferred from dialogue text. It tends to be unsurprising if there are adult scenes or drinking scenes existing in a movie without any

verbal indication.

## 6 Conclusion & Future Work

In this work, we applied a deep learning-based method to predict severity of age-restricted content based on movie script data. The experimental results show the proposed multi-task ranking-classification model outperforms the previous state-of-the-art method and can give rich interpretability by demonstrating severity using example comparator movies. Our work provides a reasonable groundbreaking exploration in this research topic for the community. For future work, we propose to investigate other modalities to capture relevant patterns and fine-grained aspects like violence types.

## Acknowledgements

This work was partially supported by the National Science Foundation under grant # 2036368. We would like to thank the anonymous EMNLP reviewers for their feedback on this work.

## References

- Craig A Anderson, Brad J Bushman, Bruce D Bartholow, Joanne Cantor, Dimitri Christakis, Sarah M Coyne, Edward Donnerstein, Jeanne Funk Brockmyer, Douglas A Gentile, C Shawn Green, et al. 2017. [Screen violence and youth behavior](#). *Pediatrics*, 140(Supplement 2):S142–S147.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Eyke Hüllermeier, Johannes Fürnkranz, Weiwei Cheng, and Klaus Brinker. 2008. [Label ranking by learning pairwise preferences](#). *Artificial Intelligence*, 172(16):1897–1916.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. [Recurrent convolutional neural networks for text classification](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. [Recurrent neural network for text classification with multi-task learning](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, page 2873–2879. AAAI Press.
- Victor Martinez, Krishna Somandepalli, Yalda Tehrani-Uhls, and Shrikanth Narayanan. 2020. [Joint estimation and analysis of risk behavior ratings in movie scripts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4780–4790, Online. Association for Computational Linguistics.
- Victor R. Martinez, Krishna Somandepalli, Karan Singla, Anil Ramakrishna, Yalda T. Uhls, and Shrikanth Narayanan. 2019. [Violence rating prediction from movie scripts](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):671–678.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Mahsa Shafaei, Adrian Pastor Lopez-Monroy, and Tamar Solorio. 2019. [Exploiting textual, visual and product features for predicting the likeability of movies](#). The 32nd International FLAIRS Conference.
- Mahsa Shafaei, Niloofar Safi Samghabadi, Sudipta Kar, and Tamar Solorio. 2020. [Age suitability rating: Predicting the MPAA rating based on movie dialogues](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1327–1335, Marseille, France. European Language Resources Association.
- Honglun Zhang, Liqiang Xiao, Yongkun Wang, and Yaohui Jin. 2017. [A generalized recurrent neural architecture for text classification with multi-task learning](#). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3385–3391.

## A Appendix

### A.1 Document Length Distribution

The document length distribution of each aspect is shown in Figure 2. Different aspects share a similar document length distribution. The average length and the median length of the movie scripts are around 10,000 words.

### A.2 Label Distribution

The severity label distribution for each aspect is unbalanced to a greater or lesser extent, as shown in Figure 3.

### A.3 Data Separation

The training, development, and test separation of each age-restricted aspect is shown in Table 4.

Aspect	Sex	Violence	Profanity	Substance	Frightening
Train	5200	3921	3910	3538	3553
Dev	651	491	489	443	445
Test	651	491	489	443	445

Table 4: The number of instances for each aspect.

### A.4 Computing Infrastructure & Settings of Experiments

We implement neural network models using PyTorch 1.6.0 and PyTorch Lightning 1.0.2. The ex-

periments are executed on NVIDIA Tesla P40.

For model development, we use Glove embedding of 300-dimension trained on Wikipedia 2014 + Gigaword 5 (6B tokens) for TextCNN, TextRCNN models. Three 2D convolution modules of TextCNN are with kernel sizes 3, 4, 5 respectively. They all have 1 input channel and 10 output channels. BERT model uses the *bert-base-uncased* model provided by the Python library *Transformers*. The proposed model use sentence embeddings of 768-dimension from SentenceTransformer based on BERT-large. Bi-LSTM used in the proposed model and TextRCNN has a hidden size of 200 for each direction. All experiments with neural networks use Adam optimizer to optimize with a learning rate set to 0.001.




Figure 2: Document length distribution for each aspect. The plots use a logarithmic scale on the x-axis (document length).




Figure 3: Label distribution for each aspect. 0-None, 1-Mild, 2-Moderate, 3-Severe.