

Automatic Characterization of Stories

by
Sudipta Kar

A dissertation submitted to the Department of Computer Science,
College of Natural Sciences and Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in Computer Science

Chair of Committee: Dr. Thamar Solorio

Committee Member: Dr. Carlos Ordonez

Committee Member: Dr. Rakesh Verma

Committee Member: Dr. Mirella Lapata

University of Houston
May 2020

Copyright 2020, Sudipta Kar

ACKNOWLEDGMENTS

April 4, 2015, 3 AM.

I just finished watching ‘The Ph.D. Movie’. I had gotten my student visa a day ago to travel to the USA and start my Ph.D. program. I still remember how scared I became after watching the movie. I was constantly thinking, should I cancel the admission and start preparing for a government job in Bangladesh or take the gamble? Five years later, I am about to finish this milestone and I am happy that I made the right decision that night. From my childhood, I am a huge fan of test cricket matches, where a single game between two teams runs for five consecutive days (each day has three sessions just like the three semesters in an academic year) and at the end, the game can produce no winning team at all. These five years of my Ph.D. journey were exactly like a test match. There were good and bad sessions, tough survival phases, dramas, moments that required strong teamwork and resilience, achievements, and proud moments. I am very happy that I had a lot of people backing me up in this period.

First and foremost, I want to thank my amazing advisor Thamar Solorio. She taught me a lot about NLP, Machine Learning, and research during this time. But the most valuable learning from her is that the most important thing to be is a good human being in the first place and have a life outside of work. She was always very supportive, kind, and inspiring. There were numerous periods of extreme frustration. Once I had a one to one meeting with her, everything was fixed within 45 minutes. I could share anything with her and I always knew that while my parents are seven thousand miles away from me, my advisor is here to take care of everything. She is such a wonderful human being. I do not know how many Ph.D. students jammed rock music with their advisors. I am happy to say that I had that chance. Thank you for everything, Professor.

I would like to thank Professor Carlos Ordonez and Rakesh Verma for serving my thesis

committee and providing helpful feedback on my work. I also thank Professor Mirella Lapata for providing tremendous support to my thesis. I had the chance of enjoying her amazing keynote speech at ACL 2017 in Vancouver, Canada. The work in this thesis is very relevant to her research and I wanted to talk to her about it. I was a second-year Ph.D. student and was very hesitant to approach such a giant of NLP. Later I wrote her an email and she surprised me by agreeing to be in my committee. Her thoughtful insights into the research positively transformed my work a lot.

I want to thank the past and current members of the RiTUAL group at UH. Everyone in this lab was super helpful, easy to collaborate. Thank you Adrián Pastor López Monroy, Amirreza Shirani, Deepthi Mave, Giovanni Molina, Gustavo Aguilar, John Arevalo, Mahsa Shafaei, Nicolas Rey Villamizar, Niloofar Safi Samghabadi, Prasha Shrestha, Shuguang Chen, Siva Uday Sampreeth Chebolu, Suraj Maharjan, Thoudam Doren Singh, and Yigeng Zhang. It was nice working with you all and I learned a lot from all of you. I want to give a special thanks to Suraj. He was a big brother in the lab who was always there helping me to learn new things, suggesting good papers, and even teaching deep learning concepts at 1 AM via hangout. I also want to thank Professor Fabio González for making us familiar with many amazing concepts of Deep Learning when he was visiting us for a semester.

NLP conferences were a happy part of this journey. I had the chance to attend many NLP conferences and interact with great minds in the community who encouraged me with a lot of ideas and suggestions. I want to thank Saif Mohammad, Rada Mihalcea, Dan Jurafsky, Ted Pederson, Preslav Nakov. Every time I met them, I came back with a lot of inspiration. I had the great opportunity to serve the NAACL Student Research Workshop' 2019 as a student chair. I want to thank Greg Durrett, Na-Rae Han, Farah Nadeem, Laura Burdick who were the co-organizers of the workshop. Organizing this workshop helped me a lot to grow as a researcher.

During my Ph.D. I had done three wonderful internships at Intel AI, GE Healthcare Data Science team, and Amazon Alexa. These three companies have very different areas of work where AI is being applied and I enjoyed a lot learning diverse things. I want to thank my amazing mentors ChinniKrishna Kothapali (Intel), Harini Eavani (Intel), Hsi-Ming Chang (GE), Ratna Saripalli (GE), Giuseppe Castellucci (Amazon), Simone Filice (Amazon), and Shervin Malmasi (Amazon).

Special thanks to Ruhul Amin Shajib, who was my thesis supervisor during my undergrad in SUST, Bangladesh. He was a great mentor throughout my bachelors and even now he did not forget his crazy student and was always there for any discussion. I want to thank Professor M. Zafar Iqbal. These are two great teachers in my life who completely transformed me and for whom today I am doing NLP research.

I am grateful to my parents, Prashad Ranjan Kar and Shiuli Bhawal for always being there. I am thankful to my wife Abantika Das Tonny, who was very supportive all the time and dealt with all of my mood swings. I want to thank my younger brother Showmick Kar for being there and indirectly pushing me to be a good example for him. My cousins, relatives, and in-laws were very supportive in this period. Thank you. I want to thank Auklesh Roy, Saimon Hossain, Nasir Uddin, Evanta Kabir, Imran Ahsan, Tahmid Rahman, and Emtiaz Ahmed for making me a part of their lives and families far away from home. Thanks to the people connected with me on social media for cheering to my every success, providing many ideas and encouragement from time to time.

While I am writing this, the whole world has been traumatized by the COVID-19 virus. Around three million people are infected by this virus, and 215K people have died. Scientists, doctors, nurses, health workers, and law enforcement officials all over the world are trying their best to help the affected people, find out a vaccine to fight this highly contagious virus, finding out important patterns from data. I want to thank them. They are the heroes.

Dear Parents,

I was in the third grade, we were going somewhere on a rickshaw when you two were telling me if I want to write books like what my favorite authors do, I need to study harder. I did not like studying but wanted to write. This little work is dedicated to you. Thank you for giving every bit of what you had for our happiness and education.

ABSTRACT

Computerized systems capable of generating high-level story descriptions have many potential real-life applications. However, enabling computers to do so requires teaching computers to obtain an abstract understanding of natural language stories algorithmically, which is one of the non-trivial problems in Artificial Intelligence and Natural Language Processing.

In this dissertation, we tackle the challenge of automatically characterizing stories at a high-level by generating a set of tags from narrative texts written in English. We start by presenting a background study on the problem, discuss the required resources for research, and propose a new corpus to facilitate research on high-level story understanding by selecting tag prediction for movies as an application of this problem. Then, we focus on designing methods for high-level story understanding from written narratives and predicting tags for movies from the written plot synopses. First, we employ a wide range of linguistic features to design a machine learning approach for generating descriptive tags for stories from narrative texts. At the next step, we design a neural methodology for modeling the flow of emotions throughout stories and enhance a system that uses a high-level representation of narrative texts to predict tags. We furthermore exploit the hierarchical structure of text documents to encode the synopses and strengthen the tag prediction mechanism. In the final part of this dissertation, we demonstrate a technique utilizing user reviews to generate tags for characterizing stories at a high-level. We made the new dataset, the source code of the systems, and a live tag prediction system publicly available to the community to encourage further exploration in the direction of automatic story characterization.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iii
ABSTRACT	vii
LIST OF TABLES	xiii
LIST OF FIGURES	xvi
1 Introduction	1
1.1 Less is more	2
1.2 Motivation: A Case Study on Movies	3
1.2.1 The Dilemma of Choice	3
1.2.2 Letting People Know before They Choose	6
1.2.3 Applicability	8
1.3 Research Objectives	9
1.4 Contributions	10
1.4.1 Publications	12
I Fundamentals of High-level Story Characterization	15
2 Background	16
2.1 Computational Narrative Studies: An Overview	16
2.1.1 Characters and their Activities	16
2.1.2 High-level Content Analysis	18
2.2 Automatic Tag Generation from Texts	19
2.3 Multi-Label Classification	20
2.4 Aspect and Opinion Extraction from Reviews	21
3 MPST: A Corpus of Movie Plot Synopses with Tags	23
3.1 Expected Properties of the Corpus	24
3.2 Towards a Fine-grained Set of Tags	25
3.3 Corpus Analysis	27
3.3.1 Multi-label Statistics	28
3.3.2 Correlation between Tags	29
3.3.3 Emotion Flow in the Synopses	30

3.4	Conclusion	32
4	Problem Formulation and Evaluation	34
4.1	Problem Formulation	34
4.2	Evaluation Principles and Challenges	35
4.2.1	Evaluating Correctness in Multi-label Classification	35
4.2.2	Evaluating Variation in Generated Tags	40
II	High-level Story Understanding	42
5	Linguistic Features	43
5.1	Feature Description	43
5.2	Experimental Setup	46
5.3	Results and Analysis	47
5.4	Conclusion	50
6	Modeling Flow of Emotions in Stories	51
6.1	Methods	52
6.1.1	CNN Based Document Representation	52
6.1.2	Encoding Temporal Emotions	54
6.2	Experiments	58
6.2.1	Data Processing and Training	58
6.2.2	Baselines	58
6.2.3	Evaluation Measures	59
6.3	Results	59
6.4	Analysis	61
6.4.1	Incompleteness in Tag Spaces	61
6.4.2	Significance of the Flow of Emotions	62
6.4.3	Learning or Copying?	64
6.4.4	Understanding Stories from Different Representations	64
6.4.5	Challenging Tags	66
6.5	Conclusion	66
7	Hierarchical Representation of Narratives	68
7.1	Model Architecture	69
7.2	Experiments	71
7.3	Results and Analysis	73
7.3.1	Quantitative Results	73
7.3.2	Inspecting the Hierarchical Representation	75
7.3.3	Inspecting Multi-label Rank	76
7.4	Conclusion	77
8	Applicability in Different Domains	79

III	Employing Reviews for Story Characterization	82
9	Retrieving Information from Reviews	83
9.1	Motivation	84
9.2	Opinion Mining vs Story Descriptor Extraction	85
9.3	Dataset	86
9.4	Modeling	86
9.5	Complementary Tag Generation from Reviews	89
9.6	Results	91
9.6.1	Human Evaluation	92
9.6.2	Information from Reviews and Synopses	95
9.6.3	Predictions for New Movies	96
9.7	Conclusion	96
10	Conclusions	98
	Bibliography	101
	Appendices	123
A	Source Code of Multi-label Rank	124
B	Human Evaluation	127
B.1	Questionnaire Interface	127
B.2	Detailed Results	131
C	Out of Domain Stories	136
C.1	Children’s Story	136
C.1.1	Cinderella	136
C.1.2	Snow White and the Seven Dwarfs	139
C.1.3	The Story of Rapunzel, A Brothers Grimm Fairy Tale	142
C.1.4	The Frog Prince: The Story of the Princess and the Frog	146
C.1.5	Aladdin and the Magic Lamp from The Arabian Nights	148
C.2	Modern Ghost Stories	154
C.2.1	The Shadows on The Wall	154
C.2.2	The Mass of Shadows	161
C.2.3	A Ghost	164
C.2.4	What Was It?	168
C.3	Literature	175
C.3.1	Romeo and Juliet	175
C.3.2	Harry Potter and The Sorcerers Stone	177
C.3.3	Oliver Twist	181
C.3.4	The Hound of the Baskervilles	183
C.4	TV Series	188

C.4.1	Game of Thrones Season 6 Episode 9: Battle of the Bastards	188
C.4.2	The Big Bang Theory Season 3 Episode 22: The Staircase Implementation	189
C.4.3	Friends Season 5 Episode 14: The One where Everybody Finds out	191
C.4.4	Narcos Season 1	192

LIST OF TABLES

3.1	Examples of movies and tags characterizing the story.	24
3.2	Brief statistics of the MPST corpus.	27
5.1	Performance of the hand-crafted features using 5-fold cross-validation on the training data. We use three metrics (<i>F1: micro averaged F1, TR: tag recall, and TL: tags learned</i>) to evaluate the features.	48
5.2	Results achieved on the test data using the best feature combination (all features) with tuned regularization parameter <i>C</i>	49
5.3	Experimental results obtained by 5-fold cross-validation using chunk-based sentiment representations. Chunk-based sentiment features were combined with the other features described in Section 5.1.	50
6.1	Performance of tag prediction systems on the test data. We report results of two setups using three matrices (TL: Tags learned, F1: Micro f1, TR: Tag recall).	60
6.2	Example of ground truth tags of movies from the test set and the generated tags for them using plot synopses and scripts.	65
6.3	Evaluation of predictions using plot synopses and scripts.	65
6.4	Percentage of the match between the sets of top five tags generated from the scripts and plot synopses.	66
7.1	Results obtained on the validation set using different methodologies on the synopses. MLR and TL stand for <i>multi-label rank</i> and <i>tags learned</i> respectively. [†] indicates fine-tuning the last layer of BERT.	74
7.2	Results obtained on the test set using different methodologies on the synopses. MLR and TL stand for <i>multi-label rank</i> and <i>tags learned</i> respectively.	75
7.3	Visualization of the Multi-label Rank metric. Each cell presents a plot of the predictions after a particular epoch. In each plot, each column represents one data instance and each row represents the rank. There are $ Y $ rows in total. The topmost cell in each column represents the class label having the highest probability score predicted by the model that makes it the top ranked class. Similarly the bottom-most cell is the lowest ranked class. If a cell represents a target class, the cell is colored as Blue	77

8.1	Generated tags from narratives that are not movie synopsis. We provide the full narrative texts in the Appendix C.	80
9.1	Statistics of the dataset. S denotes synopses and R denotes review summaries. . .	86
9.2	Results obtained on the test set using different methodologies on the synopses and after adding reviews with the synopses. MLR and TL stand for <i>multi-label rank</i> and <i>tags learned</i> respectively. *: t-test with p -value < 0.01	91
9.3	Predicted tags from our model for new movies getting released in 2019 using the plot synopses. Bold-faced tags are already found to be assigned by users in IMDb.	96
B.1	Data from the human evaluation experiment. B represents the tags predicted by the baseline system, N represents the tags predicted by our new system, and R represents the open set tags extracted from the user reviews by our system. If a tag is followed by a number in superscript, the number indicates the number of annotators who selected the tag as relevant to the story. We consider a tag as relevant if it has at least two votes. ♠ indicates the instances where our system’s predictions were more relevant compared to the baseline system, and ♣ indicates the opposite. For the rest of the instances, both systems had a tie. Annotators’ feedback about the helpfulness of the tagsets (closed set tags and open set tags) are presented by emoticons (😊: Very helpful, 😐: Moderately helpful, 😞: Not helpful). First three emoticons are the feedback from all the annotators for the tags from the baseline system and our system. Rest of the three emoticons are the feedback for the tags extracted from the user reviews.	131

LIST OF FIGURES

1.1	Wordwide movie production in 2017. Source: UNESCO Institute for Statistics. http://uis.unesco.org/en/news/cinema-data-release	4
1.2	User growth of Netflix. Source: appinventiv.com	5
1.3	Example feedback in IMDb showing the star rating and the written review provided by a user. Highlighted segments show what the particular user experienced about the movie that could be also useful for recommending this particular movie to other potential viewers.	7
1.4	Overview of the Contributions in this Dissertation.	11
2.1	Distinguishing Binary, Multi-class, and Multi-label Classification. Each circle represents a predictable class and highlighted circles represents the predicted classes.	20
3.1	Overview of the data collection process.	26
3.2	Tag cloud created by the tags from the dataset. The size of a tag depends on the tag’s frequency in the dataset.	27
3.3	Heatmap of Positive Pointwise Mutual Information (PPMI) between the tags. Dark blue squares represent high PPMI, and white squares represent low PPMI.	30
3.4	Tracking flow of emotions in the synopses of six movies. Each synopsis was divided into 10 equally sized segments based on the words and the percentages of the emotions for each segment were calculated using NRC emotion lexicons. The y axis represents the percentage of emotions in each segment, whereas, the x axis represents the segments.	31
4.1	Illustrating the challenges in evaluating multi-label classifiers using metrics like Precision, Recall, and F1 score for Top 5 predictions. Three ground truth tags (<i>a</i> , <i>b</i> , <i>c</i>) are assigned to an example instance where the number of possible tags is 12. The figure shows the predicted relevance of these tags according to seven different models (CLF-1, CLF-2,, CLF-7). ‘x’ indicates any other tag from the possible tags.	37
4.2	Illustrating the challenge in evaluating multi-label classifiers using Average Precision (AP), Mean Reciprocal Rank (MRR), and Normalized Discounted Cumulative Gain (NDCG). Dark blue boxes indicate true positives.	39

6.1	Convolutional Neural Network with Emotion Flow. The entire model is a combination of three modules. Module (a) learns feature representations from synopses using convolutional neural network. Module (b) incorporates emotion flows with module (a) to generate a combined representation of synopses. Module (c) uses these representations to predict the likelihood of each tag.	53
6.2	Tags with higher change of recall after adding the flow of emotions in CNN. . . .	62
6.3	The flow of emotions in the plots of 2 different types of movies. Each synopsis was divided into 20 segments based on the words, and the percentage of the emotions for each segment was calculated using the NRC emotion lexicons. The y axis represents the percentage of emotions in each segment, whereas, the x axis represents the segments.	63
7.1	Architecture of the hierarchical model that learns weighted representation of words and sentences in a narrative text for creating a high-level story representation.	69
7.2	Example sentences from the synopsis of the movie Rush Hour 3 with one of the most relevant tags from the sentence-level predictions. The importance of particular sentences and words for predicting tags are indicated by the highlight intensity of the sentence ids and words. Ground truth tags are <i>bleak, violence, comedy, murder</i>	76
9.1	Example snippets from plot synopsis and review of The Godfather and tags that can be generated from these.	84
9.2	At the center, we show an overview of the model that takes a plot synopsis and a review summary as inputs, uses two separate encoders for the inputs to construct the high-level representations and uses them to compute $P(Y_1)$. (a) illustrates an enhanced view of the synopsis encoder. It uses a Bi-LSTM with attention to compute a representation \mathbf{sh}_i^P for the i_{th} sentence in the synopsis. Additionally, a matrix of word-level attention weights α_w^P is generated that indicates the importance of each word in each sentence for correctly predicting $P(Y_1)$. Another attention based Bi-LSTM is used to create a synopsis representation d_h^P from the encoded sentences \mathbf{sh}^P . Additionally for each sh_i^P , sentence-level prediction $P(Y_1)_i$ is computed which is aggregated with d_h^P to create the final synopsis representation. (b) illustrates an almost similar encoder for reviews. To create an additional tagset Y_2 by mining opinions from the reviews, word-level importance scores α_w^R and sentence-level importance scores α_s^R are used). Apart from that, review representation d_h^R is computed in a similar way as in (a) which is used together with d_h^P to compute $P(Y_1)$	87
9.3	Illustration of the method for tag extraction from reviews using the importance score computed by the attention weights. All words left of the red vertical line are selected as candidate words.	90
9.4	Example sentences from the review of the movie August Rush with sentence ids and words highlighted based on importance on tag prediction. Ground truth tags are <i>thought-provoking, romantic, inspiring, flashback</i>	93

9.5	Summary of the human evaluation experiment results. (a) competing systems and the percentage of samples where their predictions were chosen as the best (b) complementary tags marked as helpful by number of annotators (c) helpfulness of the tagsets created from the closed set tagset (d) helpfulness of the complementary tagsets.	93
9.6	Percentage of gates activated for synopses and reviews. More active gates indicate more importance of the source for certain tags.	95
B.1	Questionnaire Interface Part 1: Written guideline to complete the task. . . .	128
B.2	Questionnaire Interface Part 2: Presenting a plot synopsis to the subject to read.	128
B.3	Questionnaire Interface Part 3: Evaluating the relevance and helpfulness of the closed set tags.	129
B.4	Questionnaire Interface Part 4: Evaluating the relevance and helpfulness of the open set tags extracted from the reviews.	129
B.5	Questionnaire Interface Part 5: Inspection of any possible bias in the judgement.	130
B.6	Questionnaire Interface Part 6: Collecting subjects confidence.	130

Chapter 1

Introduction

“But how could you live and have no story to tell?”

- Fyodor Dostoevsky

Stories are an inseparable part of mankind from the very beginning of the evolution of *Homo sapiens*. From the paintings on the wall of ancient caves to today’s modern civilization, it is fascinating to see how mankind kept telling their stories in many forms for centuries after centuries. From ancient times, parents tell bedtime stories to their children as an enjoyable way to teach them what is right and what is wrong. Stories help us to learn about different cultures, values, and history. We can easily take a break from our life by reading a fictional novel or watching a fantasy film. Biographical narratives of brave individuals like Martin Luther King Jr. inspire us to stand up for what we believe. Horror stories can trigger the fight-or-flight response within our body and help the secretion of feel-good hormones like Serotonin. Fairy tales make us want to believe that in the end, the good will win against all the evils. Thus, and in many other ways, stories have a large impact on our lives and in fact, we all have our own life stories and we absolutely love it when people find our stories enjoyable.

1.1 Less is more

Revolution in technology has accelerated the production of stories written for commercial purposes. Keeping pace with the growing number of consumers, more novels are being published each year and the annual film production rate is also higher than the past. Now the consumers have a lot of options to try and navigating through more options requires more time. Often people read or listen to the summary of stories to get an idea about a book or movie, which helps them to make a selection. For instance, book and movie recommendation systems like Goodreads¹ and the International Movie Database (IMDB)² offer the summary of a book or movie's storyline to help their users with making choices.

The need for summarizing the enormous amount of texts in different domains triggered the need for techniques that are capable of producing the summary of a text document. Automatic text summarization techniques are broadly divided into the following categories:

1. **Extractive Summarization:** The summary is obtained by identifying and extracting key parts of a document.
2. **Abstractive Summarization:** The summary is *generated* by understanding the semantics and topics from a document and rephrasing the document in a shorter version which is expected to be similar to how a human would rephrase the document by keeping the meaning unchanged.

Story summaries can be helpful for consumers to save the time for reading a whole book or watching a movie, but when it comes to selecting just one story to enjoy from a vast collection of options, summaries can still be too long for reading to make a selection. In such cases, a more suitable solution could be an extreme form of summarization, which is

¹goodreads.com

²imdb.com

to characterize and describe stories by a small set of short text labels called *Tags*. The advantage of tags over the summary is that tags are a much faster way to let people know the essence of a story.

Stories have life and life needs to be understood to rephrase the gist correctly. Hence, being able to automatically describe a story with a small set of tags, we need to design efficient abstraction techniques to understand a story at such a level, so that the generated tags will match what a human would use to describe the story. We define such a process as **High-level Story Understanding**. In this dissertation, we want to achieve high-level story understanding as *being able to understand the core elements of a story like a theme, emotions, events, genre, impact on consumers, and summarize these into a high-level form like tags*.

1.2 Motivation: A Case Study on Movies

1.2.1 The Dilemma of Choice

Thousands of movies are being produced every year on average (Figure 1.1). It is not surprising for anyone to go adrift in this vast ocean to find the right movie to watch. In such situations, the role of rescuers is usually taken by promotional materials, news articles, critics, and friends. With the rapid advancement of technology, numerous web-based services like IMDB, Rotten Tomatoes,³ MovieLens⁴ have been created to assist people in making their choices. In recent years, several online streaming services like Netflix⁵ and Hulu⁶ have appeared and are capable of recommending movies and TV shows to their users. Such streaming services offer thousands of movies and TV shows to watch for their users in

³<http://www.rottentomatoes.com>

⁴<http://www.movielens.org>

⁵<http://www.netflix.com>

⁶<http://www.hulu.com>

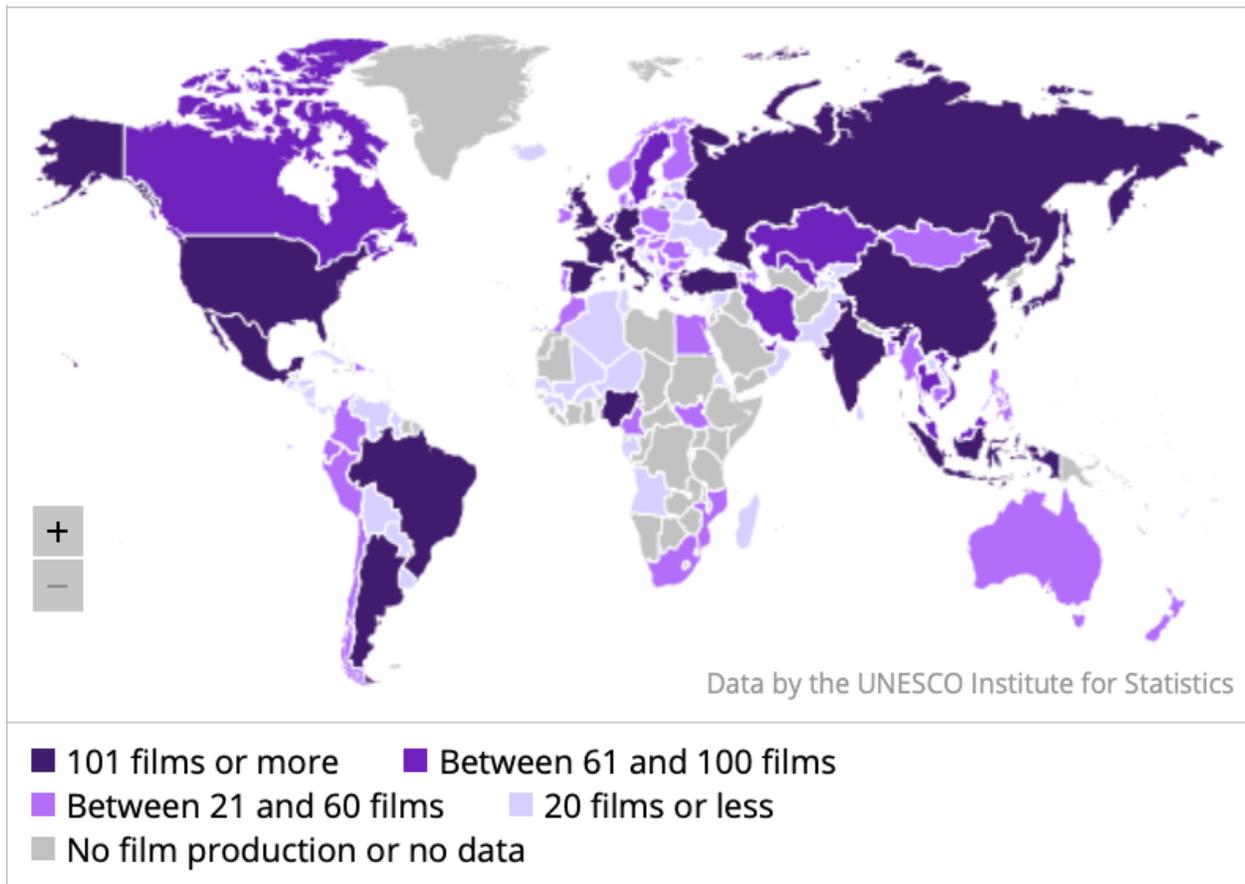


Figure 1.1: Worldwide movie production in 2017. Source: UNESCO Institute for Statistics. <http://uis.unesco.org/en/news/cinema-data-release>

exchange for a not so expensive monthly subscription fee.

Online streaming services like Netflix try to maintain a strong dynamic recommendation engine to offer personalized suggestions of movies and TV shows to their users. In general, these services recommend movies by modeling users' behavior, like what type of movies a user usually watches or what are the trending movies right now. For example, if such a system detects that a user prefers to watch a specific type of movie, it will recommend more movies of that type to that user. Such convenience attracted millions of people worldwide (Figure 1.2) to become a subscriber to these services and enjoy movies and TV shows anywhere, anytime.

Netflix Subscribers in Q3 2018

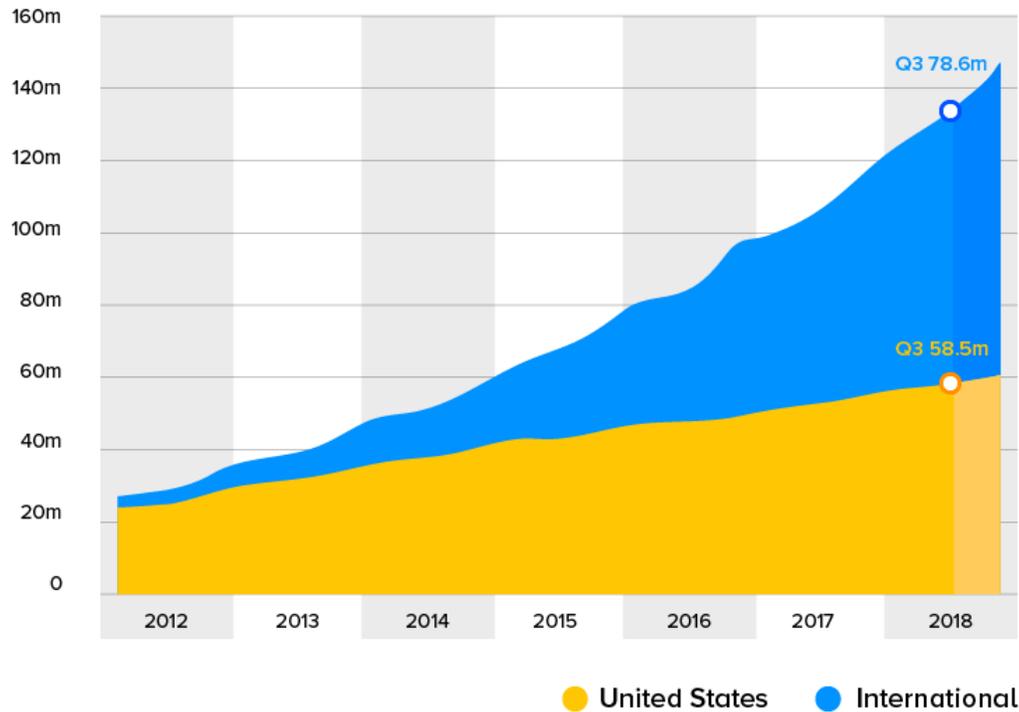


Figure 1.2: User growth of Netflix. Source: appinventiv.com

However, the current mechanism of recommendations is not always a pleasant experience for the users. Especially these systems fail to provide any clear explanation if a user wants to know, *“How is the movie? How is the storyline? Why do you think I will like this movie? I need inspiration at this moment, should I watch it?”* Because in general, such systems do not provide any effective way for the users to get the essence of any item at a glance. Even though these systems often show posters, scenes, a summary of plot, and genres to facilitate recommendations, it is a daunting task for a user to make sense of this information in a short amount of time and pick a movie to watch among the hundreds of suggested options. It leads the users to such a mental state where they do not have reasonable control over the selection process. It has been observed that possessing the control to choose something helps to feel

good and the inability to do so is naturally unpleasant and stressful [41]. Without a clear way to analyze the options, choosing a movie can be frustrating for a user and sometimes they fail to select anything at all.

1.2.2 Letting People Know before They Choose

If a system can let users know what a storyline has to offer, users will be able to gain more control over their choices. Two sources could be utilized to retrieve and present such information to users.

From the synopsis: The story is the founding block of any movie, that accompanied by audio-visual representation becomes live on screen. Major elements of stories like the characters' personality, the chemistry between the characters, and events surrounding them mostly determine what people will experience about the movie. The gist of the story could be found by reading its script or plot synopses without the need of watching the entire movie. We argue that these textual representations have sufficient information to allow users to predict what to expect from the movie.

From the people who watched it: Almost every movie recommendation service provides a mechanism for the viewers so that they can leave feedback on the movies they have watched. Star ratings and written user reviews are examples of such a mechanism. In user reviews, people can express their opinion about the movies such as, how was the story, what did they like, how they felt emotionally and so on (Figure 1.3). Such opinions often reflect story attributes and extracting them as a tagset can help other people to get an idea about the movie.

Despite the indubitable effectiveness, making sense of the story and countless reviews of a movie by reading will require a considerable amount of time for users. It is not an easy task either to grasp this large amount of texts. One way to make life easier for users is to

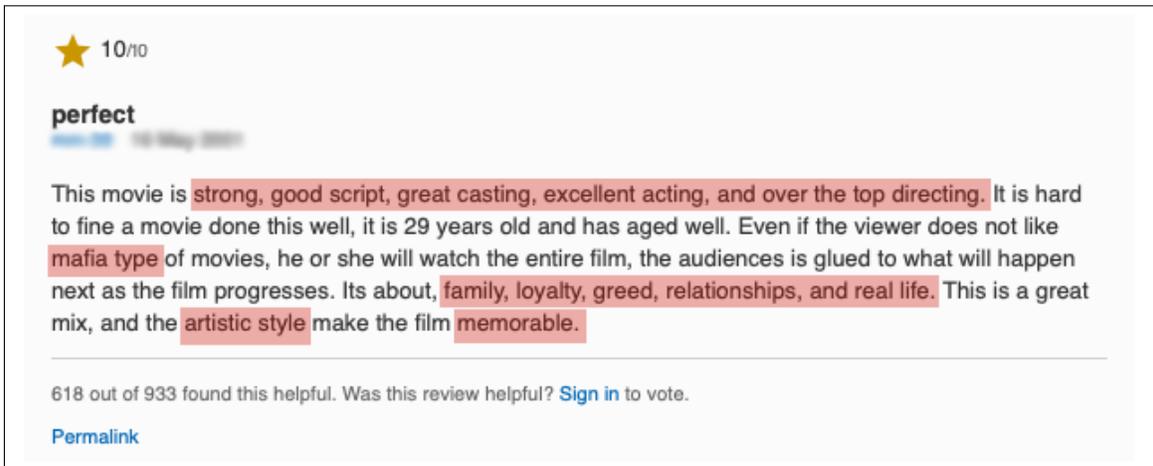


Figure 1.3: Example feedback in IMDb showing the star rating and the written review provided by a user. Highlighted segments show what the particular user experienced about the movie that could be also useful for recommending this particular movie to other potential viewers.

find out the key attributes of stories automatically from these texts and present them as tags in front of users.

User-generated tags for online items are beneficial for both the users and content providers in modern web technologies. For instance, the capability of tags in providing a quick glimpse of items can assist users to pick items precisely based on their taste and mood. On the other hand, such strength of tags enables them to act as strong search keywords and efficient features for recommendation engines [56, 97, 59, 14]. As a result, websites for different medias like photography,⁷ literature,⁸ film,⁹ and music¹⁰ have adopted this system to make information retrieval easier. Such systems are often referred to as Folksonomy [106], social tagging, or collaborative tagging.

In some movie review websites like IMDB, people can assign tags to movies after watching them. These tags often represent summarized characteristics of the movies such as emotional

⁷<http://www.flickr.com>

⁸<http://www.goodreads.com>

⁹<http://www.imdb.com>

¹⁰<http://www.last.fm>

experiences, events, genre, character types, and emotional impacts. However, this situation is not the same for all of the movies. Usually, popular movies have a lot of tags as they tend to reach a higher number of users on these sites. On the other hand, low profile movies that fail to reach such an audience have very small or empty tagsets. In an investigation, we found that $\approx 34\%$ of the movies among the top $\approx 130\text{K}$ movies of 22 genres¹¹ in IMDB do not have any tag at all. It is very likely that the lack of descriptive tags negatively affects the chances of many movies being discovered.

High-level story understanding can act as a remedy in such a situation. An automatic process to create tags for movies by high-level story understanding would reduce the dependency on humans to accumulate tags for movies and these tags can help users to make choices by providing a characterization of stories with just a few words.

1.2.3 Applicability

Although here we discuss how high-level story understanding can assist people in selecting movies, the applicability of such a method is not limited to only this domain. For example, tags generated by high-level story understanding could be employed to quickly describe story-based items in many applications like *book recommendation systems*, *community writing platforms*, *children storybooks*, *bibliotherapy*, *literary blogs*, and *news*. These tags can help people to make quicker selection of story-based items. Additionally, such tags can act as soft categories to be used in recommendation and search systems.

¹¹<http://www.imdb.com/genre/>

1.3 Research Objectives

Using a set of tags to characterize stories can be seen as a surface-level representation of stories empowered by a comprehensive dissection of the story elements. We hypothesize that such a representation of stories can be achieved from the written narratives. Based on our discussion in Section 1.2.2, we also argue that depending on the availability, user reviews can also contribute to the characterization process. Therefore, the main research question that we want to answer in this dissertation can be stated as the following:

Is it possible to develop computational approaches to extract a set of tags that provides a reasonable high-level description of a story?

In other words, we want to investigate:

1. How can we design a method for generating a set of tags representing different elements of stories through high-level story understanding?
2. How can we develop a process to identify and extract different attributes of stories from user reviews?

The first part of squeezing a story into a set of tags can be seen as an extreme form of abstraction of narrative texts, whereas the second part of detecting story attributes from the reviews can be seen as a form of extraction based method. However, our goal of abstraction of narratives is different from abstractive text summarization. While abstractive summarization on a narrative text is expected to distill major events like *boy meets girl* and *hero kills villain*, our task is to describe the entire story with high-level descriptors like *inspiring*, *though-provoking*, *violence* and so on.

We lay out the following goals that we want to address in this dissertation, which will eventually help the broader research community of Natural Language Processing (NLP):

1. **Resources:** We notice the scarcity of suitable datasets for designing methods for high-level story understanding. Thus, we aim at creating new resources as part of this dissertation to facilitate research in this direction.
2. **Computational Approaches for Story Abstraction:** We want to investigate and contribute multiple computational approaches for designing an automatic system for a high-level understanding of stories. We will experiment and analyze the applicability of a wide range of linguistic features. Additionally, we explore the feasibility of automatic representation of narrative texts produced by Artificial Neural Networks (ANN) for story characterization.
3. **Un/Weakly-supervised Story Descriptor Extraction:** We will work towards to explore the scope of story attribute extraction from user reviews. Where problems of this category are typically solved by supervised learning, we want to push this problem in such a paradigm that has no direct supervision.

1.4 Contributions

In this dissertation, we make contributions in terms of resource and methodologies (Figure 1.4) to deal with the challenge of characterizing a story with a set of tags, where we select the task of assigning tags to movies as an application of this problem. We explore two aspects of the problem as part of this work. The first part is focused on designing methodologies for high-level story understanding to infer story attributes from narrative texts, which we consider as an extreme form of abstractive summarization. In the second part, we explore the direction of extractive summarization for story characterization, where we aim at designing techniques for learning to extract story attributes from reviews without direct supervision.

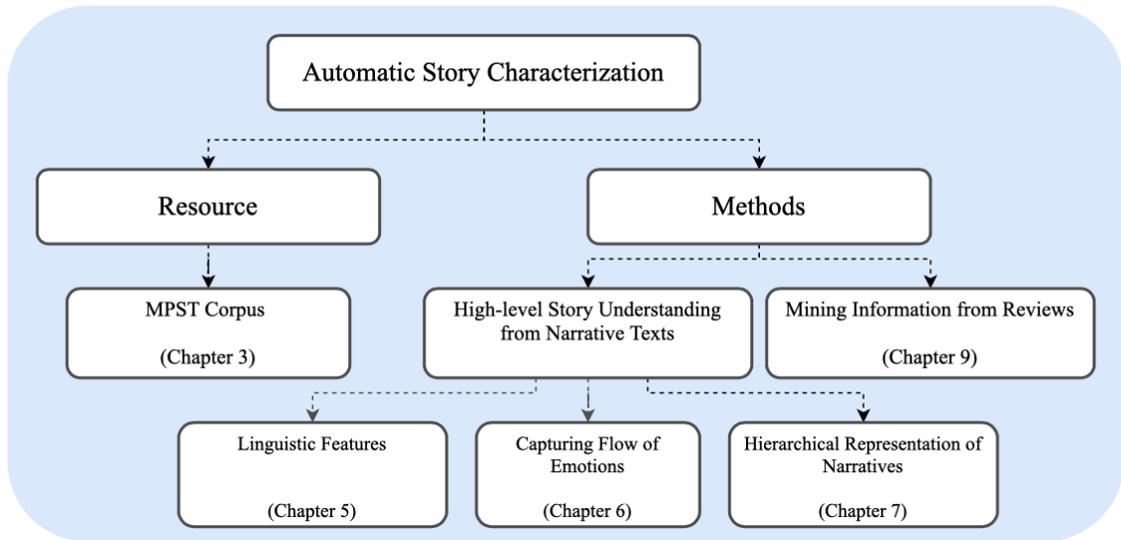


Figure 1.4: Overview of the Contributions in this Dissertation.

As part of this dissertation, we develop new resources to facilitate research on story characterization. We create a new corpus that contains the plot synopsis of $\approx 15k$ movies and a set of ≈ 70 tags to support the task of high-level story understanding. This tagset is constructed from human assigned tags on movies in different movie recommendation websites. We methodologically filter a large tag space to keep the tags that specifically describe various properties of the storyline of movies (Chapter 4). This dataset is publicly available for people to use in research purpose.^{12,13} Additionally, we compile a collection of user reviews about these movies from the web to enable research on extracting story attributes from reviews.

We build a benchmark machine learning based story representation system using a wide range of linguistic features that is capable of generating a tagset for movies given the written plot synopsis. Our experimented features capture lexical, syntactic, semantic, sentiments and emotions based properties from the narratives, where we thoroughly investigate the effectiveness of each property individually and also as part of different feature combinations. We find that the lexical representations act as strong features for such a system and chunk

¹²ritual.uh.edu/mpst-2018

¹³kaggle.com/cryptexcode/mpst-movie-plot-synopses-with-tags

based representation of sentiments and emotions increases the capability of this system to identify more attributes of stories.

At the next step, we study the scope of automatically learning the feature representations from a synopsis using Convolutional Neural Networks (CNN). Inspired by our findings of chunk-based representations of sentiments and emotions based features, we model the shifts in emotions as a story progresses using Recurrent Neural Networks (RNN). We show that combined with the textual feature representation produced by CNN, our proposed method to capture the flow of emotions results in improved learning of diverse story attributes.

Then we look beyond considering a narrative text as a plain sequence of lexical units and exploit the hierarchical representation in a document. We employ a neural model that learns to automatically create a document representation by hierarchically learning the representation and importance of the words and then the sentences in a narrative text. Our experimental results show that being able to identify the importance of each word and sentence plays a vital role in correctly identifying the tags for a story (Chapter 7).

Finally, we set out the task of extracting tags from movie reviews for story characterization. For this purpose, we design a multi-view neural model to incorporate movie reviews with the synopses to predict tags. Besides improving supervised tag prediction, such an approach enables us to extract tags from the reviews in an unsupervised fashion (Chapter 8). All of the dataset and source code produced during this work are released as public and we have observed growing interest in these resources from the NLP community.¹⁴

1.4.1 Publications

During my Ph.D. program, we worked on the following publications as part of this dissertation:

¹⁴As of February 2020, the proposed dataset has been downloaded around 700 times.

1. Sudipta Kar, Suraj Maharjan, A. Pastor López-Monroy, and Thamar Solorio. MPST: A corpus of movie plot synopses with tags. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France, May 2018. European Language Resources Association (ELRA). ISBN 978-2-9517408-9-1
2. Sudipta Kar, Suraj Maharjan, and Thamar Solorio. Folksonomication: Predicting tags for movies from plot synopses using emotion flow encoded neural network. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2879–2891, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/C18-1244>
3. Sudipta Kar, Gustavo Aguilar, and Thamar Solorio. Multi-view characterization of stories from narratives and reviews using multi-label ranking. *arXiv preprint*, 2019. URL <https://arxiv.org/pdf/1908.09083.pdf>

In addition to the mentioned papers, we published the following papers:

4. Marc Franco-Salvador, Sudipta Kar, Thamar Solorio, and Paolo Rosso. UH-PRHLT at SemEval-2016 Task 3: Combining lexical and semantic-based features for community question answering. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 814–821, San Diego, California, June 2016. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/S16-1126>
5. Sudipta Kar, Suraj Maharjan, and Thamar Solorio. RiTUAL-UH at SemEval-2017 Task 5: Sentiment analysis on financial data using neural networks. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 877–882, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/S17-2150. URL <https://www.aclweb.org/anthology/S17-2150>

6. Suraj Maharjan, Sudipta Kar, Manuel Montes, Fabio A. González, and Thamar Solorio. Letting emotions flow: Success prediction by modeling the flow of emotions in books. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 259–265, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-2042. URL <https://www.aclweb.org/anthology/N18-2042>
7. Niloofar Safi Samghabadi, Deepthi Mave, Sudipta Kar, and Thamar Solorio. RiTUAL-UH at TRAC 2018 shared task: Aggression identification. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 12–18, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/W18-4402>
8. Gustavo Aguilar, Sudipta Kar, and Thamar Solorio. LinCE: A centralized benchmark for linguistic code-switching evaluation. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, Paris, France, May 2020. European Language Resources Association (ELRA)
9. Mahsa Shafaei, Niloofar Safi Samghabadi, Sudipta Kar, and Thamar Solorio. Predicting the mpaa rating based on movie dialogues. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, Paris, France, May 2020. European Language Resources Association (ELRA)
10. Md. Zobaer Hossain, Md Ashrafur Rahman, Md Saiful Islam, and Sudipta Kar. A dataset for detecting fake news in Bangla. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, Paris, France, May 2020. European Language Resources Association (ELRA)

Part I

Fundamentals of High-level Story

Characterization

Chapter 2

Background

2.1 Computational Narrative Studies: An Overview

Computational Narrative Studies deals with developing algorithmic approaches to understand, represent, and generate natural language stories. In this area of research, computational approaches have been applied to enhance the analysis and interpretation of written narratives. We broadly divide these works into two categories, where the studies in the first category focus on the low-level analysis of the individual characters and events, and the second category of works concentrates on comprehending entire stories at a higher level.

2.1.1 Characters and their Activities

Characters are one of the key elements of a story. Therefore, studying the characters has been a significant part of computational approaches to understand and analyze stories. For instance, Propp's theory [81] influenced several works in the area of computational analysis of narratives. According to this theory, all the characters in stories can be divided into seven

abstract characters, which are *Hero*, *Villain*, *Dispatcher*, *Donor*, *(Magical) Helper*, *Sought-for-person*, and *False Hero*. Valls-Vargas et al. [104] developed a method to automatically extract the characters in folk tales and classify their structural roles according to Propp’s theory. In their work, they used *sentence segmentation*, *parts of speech tagging*, *dependency parsing*, and *co-reference resolution* to extract different character entities and verb entities. Then they identified the relationship between the characters and actions that took place in the stories. Utilizing this information, they assigned a character type to each of the characters. For identifying characters from folk tales, Declerck et al. [24] took an ontology-based approach. In a similar work, Bamman et al. [10] tried to model the latent personas and roles of film characters using a Dirichlet Persona Model that uses structured linguistic information from movie plots. Character identification is also approached in Goyal et al. [34], where the authors designed an approach to identify characters in fables. Finally, Mamede and Chaleira [67] developed the Direct or Indirect Discourse (DID) system that is capable of identifying characters in children’s stories.

Understanding the interactions between the characters in a story are sometimes helpful to understand the central theme or events of the story. For example, *conflicts* accompanied by *violence* and *negative affect* can indicate *action* or *climax* in a story. Several studies explored this challenging problem from different perspectives. For instance, Iyyer et al. [42] took an unsupervised approach to model the relationship of two characters over the time using a set of event descriptors like *marriage*, *murder*, *love*, and *sadness*. In another similar set of works, Chaturvedi et al. [19, 20] used the swing in sentiments over time to model the evolving relationship between two characters.

Identification of the characters and their interactions in stories advanced the development of systems capable of inferring the high-level interconnection between all the characters in a story through constructing social networks [1, 2, 3, 54]. Such models help to procure

the overall scenario of the interplay between all the characters in a story hence aiding the interpretation. However, as these studies focus on the low-level interpretation of stories, it is somewhat difficult to describe a story at high-level from such interpretations.

2.1.2 High-level Content Analysis

Over the years, high-level story analysis approaches primarily evolved around the problem of identifying genres [13, 50, 79, 112]. The task of genre identification deals with the development of computational methods to automatically classify a narrative text into one or more literary categories (e.g., horror, romance). One of the early approaches for this task was taken by Stamatatos et al. [96], where they first extracted various stylistic markers from texts and used those markers to classify the genre of a story. Worsham and Kalita [112] investigated how do themes evolve through the chapters of a book and how they compose the overall genre. In their work, they studied different deep learning architectures for identifying the genre of books. Battu et al. [11] used movie plot synopses to classify the genre of movies. While most of the works in this area focus on English texts, this particular work studied the problem in the context of some other languages like Hindi, Telugu, Tamil, Malayalam, Korean, French, and Japanese. Although the genre is somewhat helpful from a broad categorization perspective, in our work, we look beyond genre and focus on finding out more fine-grained characteristics of stories.

Emotional Dynamics in Stories: Emotions are part and parcel of the characters in any story. With time progressing, different events take place in stories that affect the characters. The famous American writer Kurt Vonnegut [108] once argued that “*stories can be described in terms of emotional shapes*”. Specifically, he insisted that the ups and downs of emotions in a story form a shape, which he named as Emotional Arcs. He also claimed that, there are six basic shapes of these emotional arcs, which are:

1. Rags to riches (rise)
2. Tragedy or Riches to rags (fall)
3. Man in a hole (fall-rise)
4. Icarus (rise-fall)
5. Cinderella (rise-fall-rise)
6. Oedipus (fall-rise-fall)

Later, a group of data scientists conducted an experiment on novels and found that the emotional arcs of stories are indeed dominated by six different shapes [87], as Kurt Vonnegut claimed. They concluded by providing empirical evidence that successful stories usually share similar patterns in their emotional arcs. This finding is significant for our task of high-level story understanding as unfolding the emotional arcs in computational models can benefit understanding various feeling related attributes of stories.

2.2 Automatic Tag Generation from Texts

Automatic tag generation from content-based analysis has drawn attention in different domains like music and images. For example, creating tags for music has been approached by utilizing lyrics [105, 39], acoustic features from the tracks [29, 26], categorical emotion models [51], and deep neural networks [21]. In the context of tag generation from textual content, two notable systems are AutoTag [71] and TagAssist [95], which utilized the textual content to generate tags, aggregate information from similar blog posts to compile a list of ranked tags to present to the authors of new blog posts. Similar works [48, 61, 98] focused on recommending tags to users of BibSonomy¹ (a social bookmark and publication

¹<https://www.bibsonomy.org>

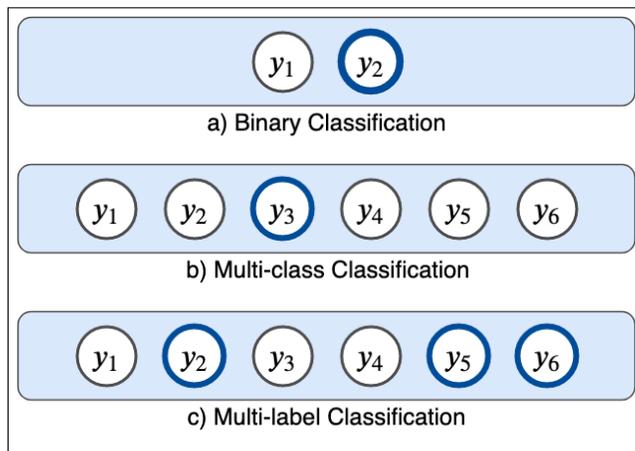


Figure 2.1: Distinguishing Binary, Multi-class, and Multi-label Classification. Each circle represents a predictable class and highlighted circles represents the predicted classes.

sharing system) upon posting a new web page or publication as proposed systems in the ECML PKDD Discovery Challenge 2008 [38] shared task. These systems made use of some kind of out of content resources like user metadata, and tags assigned to similar resources to generate tags.

In our work, we aim at looking at the textual content itself rather than using supportive metadata and making sense of the actual story. We analyze the written narratives to understand various attributes of stories and generate a set of tags to describe the stories.

2.3 Multi-Label Classification

In machine learning, classification problems can be divided into three primary categories: Binary Classification, Multi-class Classification, and Multi-label Classification. Binary classification (Figure 2.1(a)) deals with assigning a single class to a given sample of data out of two probable classes, where in Multi-class classification, the goal is to categorize a sample data into a single class out of three or more probable classes (Figure 2.1(b)). In Multi-label classification problems (Figure 2.1(c)), a sample can be categorized into one or more classes

out of the possible number of classes. Our task of predicting tags to characterize stories can be formulated as a multi-label classification problem, as for a sample story, we want to generate one or more tags.

Multi-label classification is a challenging problem in machine learning due to the complications in training models and evaluation. Such classification problems are often transformed into multiple binary classification problems [84, 85, 86], multi-class classification problems [103, 83, 101], or ensemble methods.

2.4 Aspect and Opinion Extraction from Reviews

Opinion extraction or opinion mining is a significant problem in natural language processing. The goal of an opinion extraction system is to identifying the key aspects of any product and users' opinions about those aspects. For example, an opinion mining system working on the reviews about a smartphone would try to find out what the users' are talking about key features like the screen quality, battery lifetime, design, and performance. A summarized representation of such insights (e.g., *78% of the users mentioned that the picture quality is bad compared to the price*) are extremely helpful for the future potential buyers. Moreover, this information can help the manufacturers to design their next product.

There is a subtle distinction between the reviews of typical material products (e.g., *phone, TV, furniture*) and story-based items (e.g., *literature, film, blog*). In contrast to the usual aspect based opinions (e.g., *battery, resolution, color*), reviews of story-based items often contain *end users' feelings, important events of stories, or genre related information*, which are abstract in nature (e.g., *heart-warming, slasher, melodramatic*) and do not have a very specific target aspect. Extraction of such opinions about stories has been approached by previous work using reviews of movies [118, 58] and books [60]. Such attempts are broadly

divided into two categories. The first category deals with spotting words or phrases (*excellent, fantastic, boring*) used by people to express how they felt about the story and the second category focuses on extracting important opinionated sentences from reviews and generating a summary. In our work, while the primary task is to retrieve relevant attributes from the pre-defined set of tags, we also build a system that can spot opinions.

Chapter 3

MPST: A Corpus of Movie Plot

Synopses with Tags

In the previous chapter, we shed light on different aspects related to the research on story understanding. We pointed out that high-level story understanding did not get much attention in the existing literature of natural language processing despite having the potential to be utilized in many useful applications. As a result, there is no such resource that could be directly used to build computational methods capable of extracting story related attributes from a written narrative. We identify this as a gap that needs to be filled at the beginning stage of designing systems to characterize stories for enabling research and development of story characterization systems. Therefore, we create the Movie Plot Synopses with Tags (MPST) corpus using *MovieLens 20M dataset*, *Internet Movie DataBase (IMDb)*, and *Wikipedia*. This corpus contains a fine-grained set of 71 tags that are solely related to story attributes and multi-label assignments between the tags and plot synopsis of 14,828 movies. Table 3.1 shows some samples from the MPST corpus.

Throughout the rest of this chapter, we will discuss the properties that a corpus should

Table 3.1: Examples of movies and tags characterizing the story.

	<p>A Nightmare on Elm Street 3: Dream Warriors Tags: <i>fantasy, murder, cult, violence, horror, insanity</i></p>
	<p>50 First Dates Tags: <i>comedy, prank, entertaining, romantic, flashback</i></p>

have to be useful for building story characterization systems (Section 3.1), how we developed such a corpus that meets these requirements (Section 3.2), and we provide a thorough statistical analysis on the corpus.

3.1 Expected Properties of the Corpus

As the first step towards building the MPST corpus, we defined four required properties for the corpus that must be satisfied to facilitate further research on high-level story understanding and characterization. These properties are explained below:

1. **Tags should express story attributes that are easy to understand by people.**

We identify predicting tags for movies from the plot synopses as an application of the problem we want to tackle throughout this thesis. Therefore relevant tags for movies are those that capture properties of only the stories (e.g., structure of the plot, genre, emotional experience, storytelling style), and not attributes of the movie foreign to the plot, such as metadata. More specifically, we want to infer the tags that can be directly retrieved if we have only the story as input. Additionally, we also emphasize on the fact that the tags should be easily understandable by common people. Therefore the tagset should not consist of any jargon.

2. The tagset should not be redundant.

Because we are interested in designing methods to automatically assign tags, having multiple tags that represent the same property is not desirable. For example, tags like *cult*, *cult film*, *cult movie* are closely related and should all be mapped to a single tag.

3. Tags should be well represented.

For each tag, there should be a sufficient number of plot synopses, so that the process of characterizing a tag does not become difficult for a machine learning system due to data sparseness.

4. Plot synopses should be free of noise and adequate in content.

Plot synopses should be free of noise like IMDb notifications and HTML tags. Each synopsis should have at least ten sentences, as understanding stories from very short texts would be difficult for any machine learning system.

3.2 Towards a Fine-grained Set of Tags

As shown in Figure 3.1, we collected a large number of user assigned tags from *MovieLens 20M dataset* and *IMDb*. To extract the tags commonly used by the users, we only kept the tags that were assigned to at least 100 movies. We manually examined these tags to shortlist the tags that could be relevant to the storyline of movies. We discarded the tags that do not conform to our requirements described in the previous section. At the next step, we manually examined the tags in this shortlist to group semantically similar tags together. We got 71 clusters of tags by this process and set a generalized tag label to represent the tags of each cluster. For example, *suspenseful*, *suspense*, and *tense* were grouped into a cluster labeled *suspenseful*. Through this step, we overcame the redundancy issues in the tagset and created a more generalized version of the common tags related to the plot synopses. The

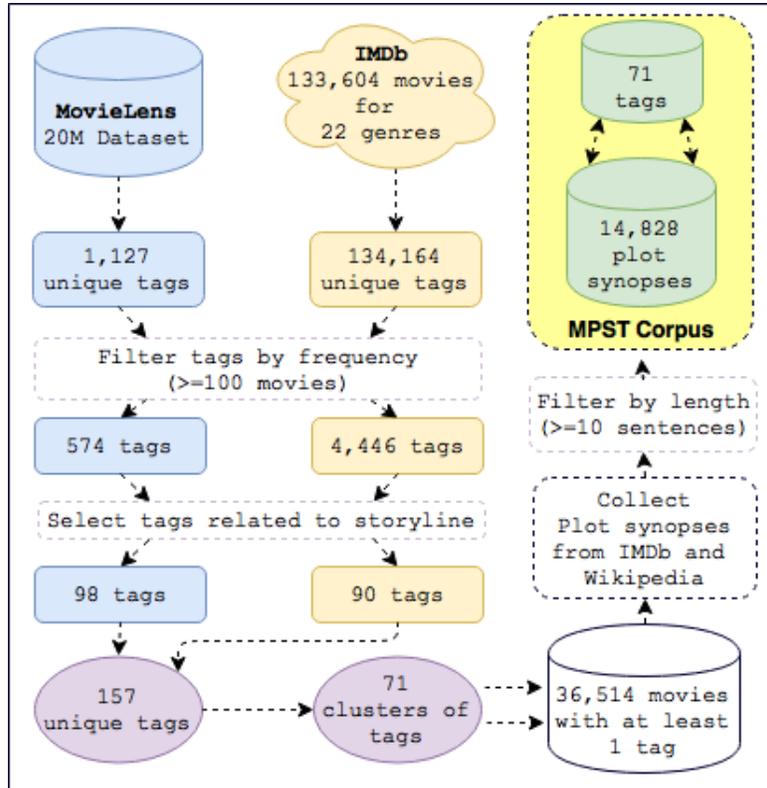


Figure 3.1: Overview of the data collection process.

tagset is shown as a word cloud in Figure 3.2.

We created the mapping between the movies and the 71 clusters using the tag assignment information we collected from *MovieLens 20M dataset* and *IMDb*. If a movie was tagged with one or more tags from any cluster, we assigned the respective cluster label to that movie. We used the IMDb IDs to crawl the plot synopses of the movies from IMDb. We collected synopses from Wikipedia for the movies without plot synopses in IMDb or if the synopses in Wikipedia were longer than the synopses in IMDb. These steps resulted in the MPST corpus that contains 14,828 movie plot synopses where each movie has one or more tags.

movie are skewed. Most of the synopses are small in terms of the number of sentences, although the corpus contains some really large synopses with more than 1K sentences. Around half of the synopses have less than 33 sentences. A similar pattern is noticeable for the average number of tags assigned to the movies. Some movies have a large number of tags, but most of the movies are tagged with one or two tags only. *Murder*, *violence*, *flashback*, and *romantic* are the most frequent four tags in the corpus that are assigned to 39%; 30%; 20% and 20% of the movies respectively. Least frequent tags like *non-fiction*, *christian film*, *autobiographical*, and *suicidal* are assigned to less than 55 movies each.

3.3.1 Multi-label Statistics

Label cardinality (LC) and *label density* (LD) are two statistics that can influence the performance of multi-label learning methods [100, 102]. Label cardinality is the average number of labels per example in the dataset as defined by Equation 3.1.

$$LC(D) = \frac{1}{|D|} \sum_{i=1}^{|D|} |Y_i| \quad (3.1)$$

Here, $|D|$ is the number of examples in dataset D and Y_i is number of labels for the i^{th} example. Label density is the average number of labels per example in the dataset divided by the total number of labels, as defined by Equation 3.2.

$$LD(D) = \frac{1}{|D|} \sum_{i=1}^{|D|} \frac{|Y_i|}{|L|} \quad (3.2)$$

Here, $|L|$ is the total number of labels in the dataset. [12] analyzed the effects of cardinality and density on multiple datasets. They showed that, for two datasets with similar cardinalities, learning is harder for the one with lower density. And if the density is similar, learning is harder for the one with higher cardinality. For example, learning performance was

better for the *Genbase* dataset (LC: 1.252, LD: 0.046) as compared to the *Medical* dataset (LC: 1.245, LD: 0.028), where they had similar cardinalities but the *Medical* dataset was less dense. On the other hand, performance was better for the *Emotions* dataset (LC: 1.869, LD: 0.311) as compared to the *Yeast* dataset (LC: 4.237, LD: 0.303), where they had similar density but cardinality of the *Yeast* dataset was higher. The label cardinality and label density of our dataset are *2.98* and *0.042*, respectively. Based on the mentioned experiments, we suspect that a traditional multi-label classification approach for this dataset will be a challenge that opens the scope for exploring more scalable approaches.

3.3.2 Correlation between Tags

To find out significant correlations in the tagset, we compute the Positive Pointwise Mutual Information (PPMI) between the tags, which is a modification over the standard PMI [22, 23, 77]. PPMI between two tags t_1 and t_2 is computed by the following equation:

$$PPMI(t_1; t_2) \equiv \max(\log_2 \frac{P(t_1, t_2)}{P(t_1)P(t_2)}, 0) \quad (3.3)$$

where, $P(t_1, t_2)$ is the probability of tags t_1 and t_2 occurring together and $P(t_1)$ and $P(t_2)$ are the probabilities of tag t_1 and t_2 , respectively. Figure 3.3 shows the heatmap correlation of PPMI values between a subset of tags. The figure shows interesting relations between the tags and supports our understanding of the real world scenario.

High PPMI scores show that *cute*, *entertaining*, *dramatic*, and *sentimental* movies can evoke *feel-good* mood, whereas lower PPMI scores between *feel-good* and *sadist*, *cruelty*, *insanity*, and *violence* suggest that these movies usually create a different type of impression on people. Also note that, these movies have stronger relations with *horror*, *cruelty*, and *darkness*, which are not typically co-related with *feel-good* mood. These observations suggest

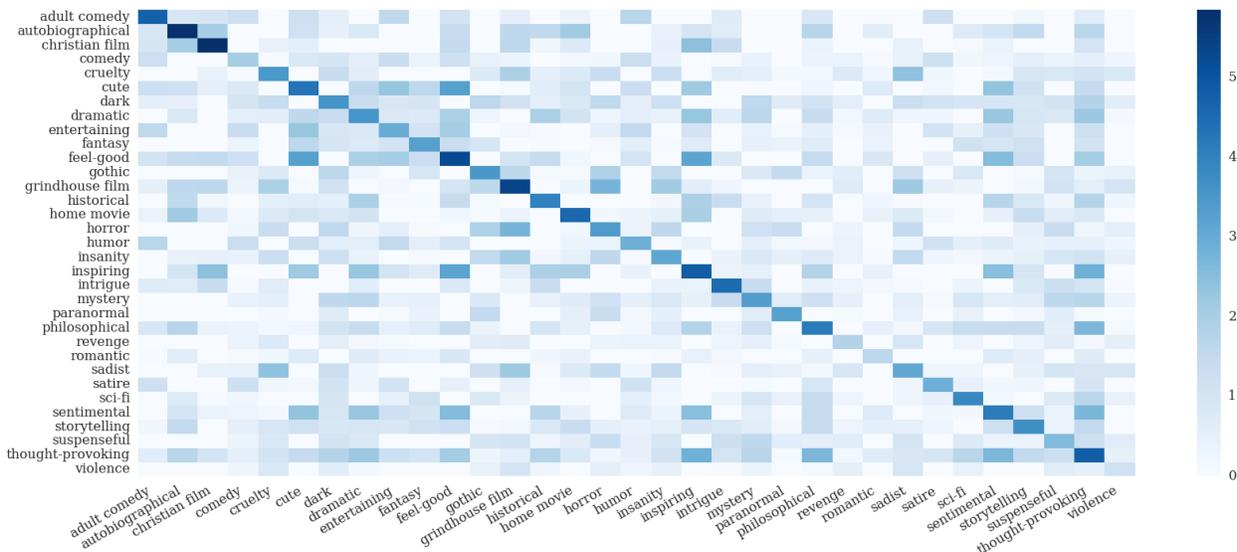


Figure 3.3: Heatmap of Positive Pointwise Mutual Information (PPMI) between the tags. Dark blue squares represent high PPMI, and white squares represent low PPMI.

that people tend to get *inspiration* from *dramatic*, *thought-provoking*, *historical*, and *home movies*. *Christian films* and *science fictions* are also good sources of *inspiration*. *Grindhouse*, *Christian*, and *non-fiction* films do not usually have *romantic* elements. *Romantic* movies are usually *cute* and *sentimental*. *Autobiographical* movies usually have *storytelling* style and they are *thought-provoking* and *philosophical*. These relations, in fact, show that the movie tags within our corpus seem to portray a reasonable view of movie types based on our understanding of possible impressions from different types of movies.

3.3.3 Emotion Flow in the Synopses

NRC Emotion Lexicons [74] have been shown effective to capture the flow of emotions in narrative stories [73]. It is a list of 14,182 words¹ and their binary associations with eight types of elementary emotions from the Hourglass of Emotions model [16] (*anger*, *anticipation*, *joy*, *trust*, *disgust*, *sadness*, *surprise*, and *fear*) with polarity. In Figure 3.4, we try to

¹Version 0.92

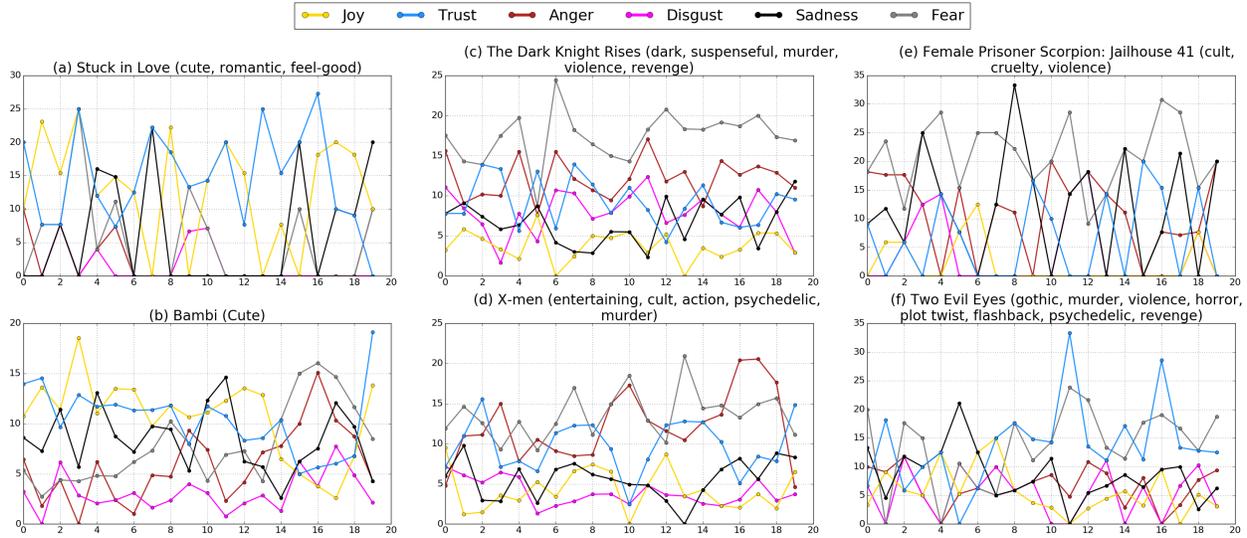


Figure 3.4: Tracking flow of emotions in the synopses of six movies. Each synopsis was divided into 10 equally sized segments based on the words and the percentages of the emotions for each segment were calculated using NRC emotion lexicons. The y axis represents the percentage of emotions in each segment, whereas, the x axis represents the segments.

inspect how the flows of emotions look like in different types of plots. The reason behind this investigation is to get a shallow idea about the potential feasibility of the collected plot synopses to predict tags.

This analysis will help us to verify if the emotional dynamics found from the synopses align with the tags. For example, in a real world scenario we will expect that horror movies will contain fear and sadness. On the other hand, comedy or funny movies will be filled with happiness.

In Figure 3.4, we can observe that, emotions like *joy* and *trust* are dominant over *disgust* and *anger* in *cute*, *feel-good*, and *romantic* movie’s plots (a, b). We can observe sudden spikes in *sadness* in segment 4. The animated movie *Bambi* (1942) shows an interesting flow of different types of emotions. The dominance of *joy* and *trust* suddenly gets low at segment 14 and gets high again at segment 18, where *fear*, *sadness*, and *anger* get high at segment 14. It is quite self-explanatory that the plot are mixtures of positive and negative emotions where the lead characters go through difficult situations, fight enemies and face

a happy ending (spike in *joy* and *trust* at the end) after climax scenes where enemies get defeated. The final segments of (b) indicate happy endings, but the rise of *sadness* and *fear* in (a) indicates that *Stuck in Love (2012)* does not have a happy ending.

We observe the opposite scenarios in cases of *violent*, *dark*, *gothic*, and *suspenseful* movies (c, d, e, and f) where *fear*, *anger*, and *sadness* dominate over *joy* and *trust*. The dominance of *anger* and *fear* is a good indicator of a movie having *action*, *violence*, and *suspense*. *Female Prisoner Scorpion: Jailhouse 41 (1972)* (e), has dominance of *fear*, *sadness*, and *anger* throughout the whole movie, and it is easy to guess that this movie has *violence* and *cruelty* portrayed through the lead characters. The flow of *joy*, *trust*, *sadness*, and *fear* alters at the middle of the movie *Two Evil Eyes (1990)* (f). Maybe it is the reason why people tagged it with *plot twist*. These observations give evidence of the connection between the flow of emotion in the plot synopses and the experience people can have from the movies, and they also match with what we expected.

3.4 Conclusion

In this chapter we presented a new corpus to contribute new resources to accelerate research on high-level story understanding. The proposed MPST corpus will be helpful to analyze and understand the linguistic characteristics of short stories like the plot synopses of movies, which will in turn help to model certain types of abstractions as tags. For example, what type of events, word choices, character personas, relationships between characters, and plot structure make a story *mysterious* or *suspenseful* or *paranormal*? Such investigations can help the research community to better exploit high-level information from narrative texts, and also help to build automatic systems to create tags for stories. We expect that the methodologies designed using this corpus could also be used to perform high-level story

understanding in other domains, such as books and storyline of video games.

Chapter 4

Problem Formulation and Evaluation

In this chapter, we will focus on the technical aspects of the problem of high-level characterization of stories. We will shed light on the task formulation, prediction strategies, and discuss about the evaluation approaches.

4.1 Problem Formulation

Let's assume, we have $|X|$ samples in a dataset, where each sample is associated with one or more tags¹ from a tagset Y . We want to learn a multi-label classification model f , which will maximize $P(Y|X)$. Therefore, given a sample x , the model will produce a probability distribution $P(Y)$ over the tagset Y .

There are two ways to generate a tagset y_{pred} ($1 \geq |y_{pred}| \leq |Y|$) for x from $P(Y)$. The first approach is to set a threshold (typically 0.5) and select all the tags with predicted probability above this threshold. The issue with this approach is that the number of predicted tags for the samples can vary a lot and for many examples only a single tag or none could be selected, which is not helpful at all to describe a story. For this reason, we follow a second route,

¹We will use the terms *tag* and *label* interchangeably in this chapter.

where we select the most probable N tags for the given instances based on $P(Y)$, hence, $|y_{pred}| = N$ for all the samples, where $N \in \mathbb{Z}^+ : 1 \leq N \leq |Y|$. In this work, we primarily experiment with $N = 3$ as the average number of tags per movie is around three (Section 3.3). Additionally, we report results with higher values of N such as 5, as more tags can provide more information about a story.

4.2 Evaluation Principles and Challenges

For a given model f generating a tagset y_{pred} for an instance x , we want to evaluate the following aspects:

1. **Correctness:** How many of the predicted tags y_{pred} align with the ground truth tagset y_{true} ?
2. **Variation:** Is the model predicting only a small portion of tags from Y or it has learned to predict all or most of the tags? If the tagset $|Y|$ is imbalanced (some tags are more frequently assigned to instances than the others), it is likely to happen that the top N predictions by a model f always have those most frequent tags and the model never learns the less frequent tags. Although the most frequent tags can be relevant, we want to make sure the model also learns to rank the less frequent tags at the top, when appropriate.

4.2.1 Evaluating Correctness in Multi-label Classification

Wu and Zhou [113] illustrate the complications in evaluating multi-label classifiers by an example of determining the significance of mistakes for the following cases: *one instance with three incorrect labels* vs. *three instances each with one incorrect label*. It is complicated to tell which of these mistakes is more serious. Due to such complications, several evaluation

methodologies have been proposed for this type of tasks [100, 113]. For example, *hamming loss*, *average precision*, *ranking loss*, *one-error* [90, 32], *micro* and *macro* averaged versions of *F1* and *AUC* score [102, 103, 62]. We found that when the average number of labels per sample is very small compared to the size of $|Y|$ (e.g., 4% in our problem), interpreting the results can be challenging with metrics like *hamming loss*, *ranking loss*, and *one-error*. As the distribution of the tags in our dataset is skewed (Section 3.3), we plan to use micro averaged F1 (micro-F1) over macro averaged F1 (macro-F1) to measure the correctness of the models in this work. Micro-F1 is computed from the micro averaged precision (P) and micro averaged recall (R) using the following equation:

$$F1 = \frac{2 \times P \times R}{P + R} \quad (4.1)$$

Where, micro averaged precision (P) and micro averaged recall (R) for $|Y|$ tags are computed using the following equations:

$$P = \frac{\sum_{i=1}^{|Y|} TP_i}{\sum_{i=1}^{|Y|} TP_i + \sum_{i=1}^{|Y|} FP_i} \quad (4.2)$$

$$R = \frac{\sum_{i=1}^{|Y|} TP_i}{\sum_{i=1}^{|Y|} TP_i + \sum_{i=1}^{|Y|} FN_i} \quad (4.3)$$

Here, TP, FP, and FN stand for *true positive*, *false positive*, and *false negative*, respectively.

4.2.1.1 Evaluation Challenges

As we have mentioned in the previous sections, for a given sample story x , we will use a model to compute $P(Y|x)$, and select the top N (3 or 5) predictions to generate a tagset to characterize that sample. The number of ground truth tags for a sample is not fixed as it depends on how many tags were assigned to the movie by the users. Hence, the number of

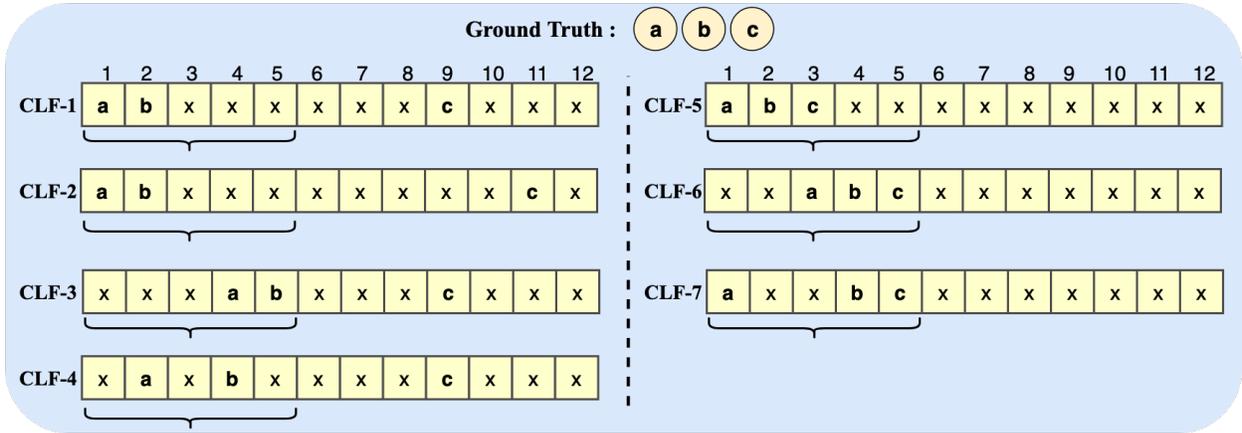


Figure 4.1: Illustrating the challenges in evaluating multi-label classifiers using metrics like Precision, Recall, and F1 score for Top 5 predictions. Three ground truth tags (a , b , c) are assigned to an example instance where the number of possible tags is 12. The figure shows the predicted relevance of these tags according to seven different models (CLF-1, CLF-2, ..., CLF-7). ‘x’ indicates any other tag from the possible tags.

target tags can be greater or lesser than our chosen N . For this reason, evaluating the top N predicted tags can hurt the evaluation for those samples if we measure the performance with precision, recall, and F1 score. For example, if a sample has two target tags assigned to it and we assign the most probable five tags according to the model to that sample (including the target tags), three of the predicted tags will be considered as false positives (FP). However, these additional three tags can be still relevant to the story while not being included in the ground truth tags because of the incompleteness problem [48] observed in social tagging platforms. Considering the reverse case scenario, where a sample has for example 15 ground truth tags, assigning five correct tags can still penalize the model as there are 10 false negatives (FN).

Another evaluation challenge is that the metrics like precision, recall, and F1 score for the top N predictions cannot measure the capability of multiple competing models in ranking the most relevant tags. We illustrate such example scenarios in Figure 4.1, where the number of tags in a classification problem is 12 ($|Y| = 12$) and we select the top 5 tags from

each classifier based on the predicted probability. Let's assume that an example instance x has three ground truth tags (a , b , c) and we are evaluating the performance of seven competing models. Two of the three ground truth tags (a and b) have appeared at the top 5 predictions for the four competing models (CLF-1, CLF-2, CLF-3, CLF-4) at the left. Hence, precision@5, recall@5, and F1@5 scores are exactly the same for all of these four models. However, CLF-1 performed better than CLF-2 in terms of predicting the relevance of c . Even though CLF-3 and CLF-4 has a and b in the top five predictions, the overall rank (4 and 5 by CLF-3, 2 and 4 by CLF-4) of these tags is not as good compared to CLF-1 and CLF-2 (1 and 2). The models at the right (CLF-5, CLF-6, CLF-7) have all the three ground truth tags in the top five predictions. But in terms of ranking the most relevant tags at the top positions, CLF-5 is better than CLF-6 and CLF-7. This aspect is not captured properly by precision, recall, and F1, which limits the comparison between multiple models.

A possible approach of addressing such ranking aspect is using rank focused metrics like Average Precision (AP), Mean Reciprocal Rank (MRR), or Normalized Discounted Cumulative Gain (NDCG). However, these metrics also have some limitations to be applied in this work. As an example, we illustrate the challenge of evaluation with these existing metrics in Figure 4.2, where we show the rank of ground truth tags predicted by two competing model CLF-1 and CLF-2. In CLF-1, all the ground truth tags rank near the top (3, 5, 6) and AP is 0.46. In CLF-2, two of the tags are ranked at top five and the remaining one near the bottom, but the AP is 0.57, which is higher than CLF-1. We observe similar results with MRR (0.29 vs 0.43) and NDCG (1.37 vs 1.68). If we apply the top N scheme to select five tags from these two models, CLF-1 will have two ground truth tags and CLF-2 will have one ground truth tag despite having a higher score in terms of AP, MRR, and NDCG. Hence, this particular metric favors the model that can rank one or two ground truth tags at the top, but not the one that does a good job at ranking all the ground truth tags near the top.

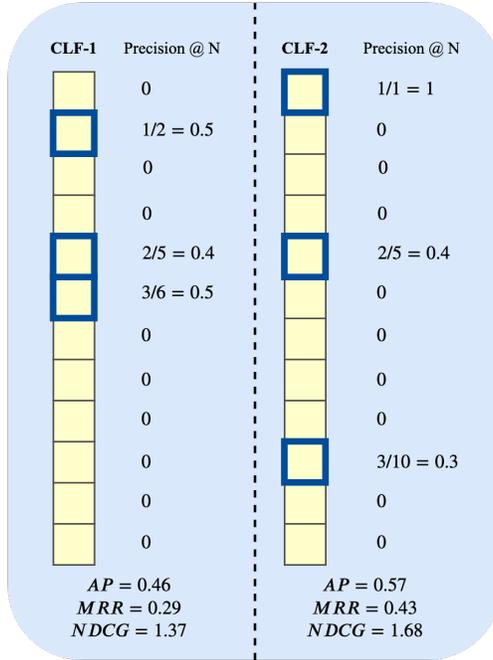


Figure 4.2: Illustrating the challenge in evaluating multi-label classifiers using Average Precision (AP), Mean Reciprocal Rank (MRR), and Normalized Discounted Cumulative Gain (NDCG). Dark blue boxes indicate true positives.

But we want to be able to choose such a model that is capable of ranking the most relevant tags at the top, so that we can confidently select the top N predictions. With this intention, we propose a new metric called Multi-label Rank.

4.2.1.2 Multi-label Rank

To overcome the evaluation challenges discussed in the previous section, we propose a new metric called Multi-Label Rank (MLR) in this work. MLR is designed to favor the model that can rank all the ground truth tags near the top like CLF-1 in Figure 4.2. MLR of D data samples is defined by

$$mlr = \frac{\sum_{i=1}^D (1 - \delta_i)}{D} \quad (4.4)$$

Here, δ_i is the mean *Rank Distance* for an instance X_i . *Rank Distance* δ_{ij} for a ground truth tag $t_{ij} \in t_i$ is defined by:

$$\delta_{ij} = \frac{\max(0, \text{rank}(t_{ij}) - |t_i| + 1)}{|T| - |t_i|} \quad (4.5)$$

where, $\text{rank}(t_{ij})$ is the ranking position of tag t_{ij} in the predicted ranking of the label set and $|T|$ is the size of label set in the dataset. This metric is helpful as it considers all the target tags as equally weighted and results in rank score of 100 if all of the K target tags are located at the top K positions in the predicted ranking regardless of the order of them. So the objective for a model becomes to have all the target tags at the top of the predicted ranking in any order to achieve a rank score of 100. For example, MLR for CLF-1 and CLF-2 in Figure 4.2 is 81.5 and 66.7 respectively. The source code of MLR is provided in Appendix A.

4.2.2 Evaluating Variation in Generated Tags

As we have discussed earlier, less frequent tags could be under-represented by models, but an ideal model should be able to discriminate among all the possible labels. Such an issue is very common in problems like image annotation, and existing works use *mean per label recall* and *labels with recall > 0* to measure the effectiveness of models in learning individual labels [57, 30, 18, 110]. Here, we use two similar metrics: *tags learned* (TL) and *tag recall* (TR). Tags learned (TL) computes how many unique tags (tags with *recall* > 0) are being predicted by the system for the evaluation set. Between two models having similar F1 score, the one that has higher TL is better as it can diversify the generated tagsets.

Tag recall (TR) computes the average recall per tag and is defined by the following equation.

$$TR = \frac{\sum_{i=1}^{|Y|} |R_i|}{|Y|} \quad (4.6)$$

Here, $|Y|$ is the size of tagset in the corpus, and R_i is recall of i^{th} tag. These evaluation metrics will help us to investigate how well and how many distinct tags are being learned by the models.

Part II

High-level Story Understanding

Chapter 5

Linguistic Features

Our first methodological contribution to high-level story understanding is a benchmark system that uses a variety of traditional linguistic features. We start exploring the problem of high-level story understanding by developing a machine learning model using this approach because such linguistic features are typically simple, easy to use, and provide vital insights about any problem while helping to achieve good performance in different NLP problems. For building a linguistic feature-based narrative representation to train a machine learning model, we experiment with capturing different types of lexical, semantic, and sentiment related characteristics from the texts. In this chapter, we will first discuss the motivation and methods of using these linguistic features. Then we will describe our experimental methodologies to representing a synopsis with the features and training a machine learning model to predict tags for movies.

5.1 Feature Description

1. **Lexical:** Representing a text document as a bag of words is a popular and effective method in various text classification problems. Hence, we adopt this model in our problem

by capturing different lexical properties from the synopses and use them as a bag. We extract word n -grams ($n=1,2,3$), character n -grams ($n=3,4$) and two skip n -grams ($n=2,3$) from the plot synopses as they are widely accepted as effective lexical representations. We use term frequency-inverse document frequency (TF-IDF) as the weighting scheme.

2. Sentiments and Emotions: Sentiments and emotions are inherent part of stories and one of the key elements that determine the possible experiences found from a story. For example, *depressive* stories are expected to be full of *sadness*, *anger*, *disgust* and negativity, whereas a funny movie is possibly full of *joy* and *surprise*. In this work, we employ two approaches to capture sentiment and emotions related features.

- **Bag of Concepts:** As concept-level information have showed effectiveness in sentiment analysis [15], we extract around 10K unique concepts from the plot synopses using the Sentic Concept parser.¹ This parser breaks sentences into verb and noun clauses and extracts concepts from them using bigram rules [82] based on the Parts of Speech (POS). For a synopsis, we extract the concepts using this concept parser and represent the synopsis as a bag of concepts in a similar way that we follow for the lexical features.
- **Affective Dimensions Scores:** The hourglass of emotions model proposed by Cambria et al. [16] categorized human emotions into four affective dimensions (*attention*, *sensitivity*, *aptitude*, and *pleasantness*) starting from the study on human emotions by Plutchik [80]. Each of these affective dimensions is represented by six different activation levels called *Sentic levels*. These make up to 24 distinct labels called *elementary emotions* that represent the total emotional state of the human mind. SenticNet 4.0 [17] knowledge base consists of 50,000 common-sense concepts with their semantics, polarity value and scores for the basic four affective dimensions. We used this knowledge

¹<https://github.com/SenticNet/concept-parser>

base to compute the average polarity, attention, sensitivity, aptitude, and pleasantness for a synopsis and use them as numeric features.

For using the sentiment and emotion features, we experiment with creating a unified feature representation for an entire synopsis and also dividing a synopsis into multiple chunks and creating feature representation for each chunk. More specifically, we divide the plot synopses into three equal chunks based on words and extracted these two sentiment features for each chunk. We will discuss more on the chunk-based sentiment representation later in Section 5.3.

3. Semantic Roles: Semantic Role Labeling (SRL) is a useful technique to assign abstract roles to the arguments of predicates or verbs of sentences. Such abstract roles are often helpful to represent the semantics in a text. We use the SEMAFOR² frame-semantic parser to parse the frame-semantic structure using the FrameNet [9] frames. For each synopsis, we use the bag of frames representation weighted by normalized frequency as feature.

4. Distributed Representation of Words: Distributed word representations or word embeddings [70] have shown effectiveness in text classification problems by capturing semantic information. Hence, in order to capture the semantic representation of the plots, we average the word vectors of every word in the plot. We use the publicly available FastText pre-trained word embeddings³ as these embeddings use the subword information to compute a representation of a word, which have been shown benefits in various text understanding tasks.

5. Agent Verbs and Patient Verbs: Actions done and received by the characters in a story can help understanding the type of the story. For example, if the characters of a story *kill, take revenge, shoot, smuggle, chase*; we can expect violence, murder, action from that story. Extracting the agent and patient verbs from the texts have been shown to be effective

²<http://www.cs.cmu.edu/~ark/SEMAFOR>

³<https://github.com/facebookresearch/fastText/blob/master/pretrained-vectors.md>

to capture such actions [10]. Hence, we use the agent and patient verbs found in synopses to extract the actions of the characters in the stories. In this regard, we use the Stanford CoreNLP library to parse the dependencies of the synopses. Then we extract the agent verbs (using *nsubj* or *agent* dependencies) and the patient verbs (using *dobj*, *nsubjpass*, *iobj* dependencies) as described in [10]. We group these verbs into 500 clusters using the pre-trained word embeddings with the K-means clustering algorithm [64] to reduce noise. We use the distribution of these clusters of the agent verbs and patient verbs over the synopses. We experimented with different values of K (K=100, 500, 1000, 1500), and 500 clusters helped to achieve better results.

5.2 Experimental Setup

Section 3.3 shows that the distribution of the number tags assigned to per movies is skewed. The average number of tags per movie is approximately three. We thus begin by experimenting with predicting a fixed number of three tags for each movie. Moreover, to get more detailed idea about movies, we create another set of five tags by predicting two additional tags.

We use random stratified split to divide the data into 80:20 train to test ratio. We use the One-versus-Rest approach to predict multiple tags for an instance. In this approach, separate binary classifiers are trained for each class and their predictions are used to make the final decision. We experiment with logistic regression as the base classifier. We run five-fold cross-validation on the training data to evaluate different features and combinations. We tune the regularization parameter (C) using a grid search technique over the best feature combination that includes all of the extracted features. We use the best parameter value ($C=0.1$) for training a model with all the training data and used that model for predicting tags for the test data.

Majority and Random Baseline: We define majority and random baselines to compare the performance of our proposed model in the task of predicting tags for movies. The majority baseline method assigns the most frequent three or five tags to all the movies. We chose three tags per movie as this is the average number of tags per movie in the dataset. Similarly, the random baseline assigns at random three or five tags to each movie.

Evaluation Metrics: As discussed in Chapter 4, we aim at evaluating two aspects of the tag prediction systems, which are *correctness of the predictions* and *variation in the generated tags*. In this chapter, we use micro-F1 for measuring the correctness and tag recall (TR) and tags learned (TL) for evaluating the variation in the generated tags.

5.3 Results and Analysis

Table 5.1 shows the performance of the hand-crafted features for predicting tags for movies. All the features beat the baselines in terms of micro-F1 and tag recall (TR). But another significant criterion to evaluate the performances is the number of unique tags predicted by the models, which is measured by the tags learned (TL) metric. We prefer such a model that is capable of creating diverse tagsets by capturing varieties of attributes of movies with reasonable accuracy. For instance, the random baseline used all of the tags in the dataset to assign to the movies but its accuracy is very poor. On the other hand, the majority baseline has better accuracy but it does not have diversity in the tagset. We can see that most of the individual features achieved almost similar micro-F1 scores, but they demonstrate difference in effectiveness to create diversity in predicted tags. Feature combinations seem to improve in TR and TL, but micro-F1 scores are almost similar to the individual features.

The lexical features show better performance compared to other features. Bag of concepts (BoC) shows similarity in performance. Combination of all lexical features demonstrates effectiveness in capturing a wide range of attributes of movies from the synopses, which is

Table 5.1: Performance of the hand-crafted features using 5-fold cross-validation on the training data. We use three metrics (*F1: micro averaged F1, TR: tag recall, and TL: tags learned*) to evaluate the features.

	Top 3			Top 5		
	F1	TR	TL	F1	TR	TL
Baseline: Most Frequent	29.7	4.225	3	31.5	7.042	5
Baseline: Random	4.20	4.328	71	5.40	7.281	71
Unigram (U)	37.6	7.883	22.6	37.1	11.945	27.4
Bigram (B)	36.5	7.216	19.6	36.1	10.808	24.8
Trigram (T)	31.3	5.204	15.4	32.4	8.461	21
Char 3-gram (C3)	37.0	7.419	22	36.6	11.264	27.4
Char 4-gram (C4)	37.7	7.799	22.6	37.0	11.582	27.2
2 skip 2 gram (2S2)	34.2	6.289	19.4	34.5	9.875	25.2
2 skip 3 gram (2S3)	30.8	4.951	12.8	32.1	8.109	18.2
Bag of Concepts (BoC)	35.7	7.984	29	35.9	12.473	34.8
Concepts Scores (CS)	31.1	4.662	7.8	32.4	7.512	8.2
Word Embeddings	36.8	6.744	13.2	36.1	10.074	17.8
Semantic Frame	33.4	5.551	13.4	33.9	8.394	15.2
Agent Verbs	32.9	5.050	7.2	33.2	7.714	8
Patient Verbs	33.1	5.134	7.4	33.5	7.843	8
U+B+T	37.2	8.732	30	36.8	13.576	36.8
C3+C4	37.8	8.662	28.8	37.4	13.395	33.6
U+B+T+C3+C4	37.1	9.991	36.8	36.8	15.871	45.8
All lexical	36.7	10.046	37.6	36.5	15.838	46.4
BoC + CS	35.7	8.165	29.4	36.0	12.754	35.4
All features	36.9	10.364	39.6	36.8	16.271	47.8

reflected by the better TR and TL scores.

We present the results achieved on the test data in Table 5.2. Although the result is similar to the result we got with all features during cross-validation, the number of predicted unique tags is higher in the test set. This result could be used as a baseline system to compare other methods developed in the future as it uses the combination of several traditional linguistic features to predict tags.

Table 5.2: Results achieved on the test data using the best feature combination (all features) with tuned regularization parameter C .

	Top 3			Top 5		
	F1	TR	TL	F1	TR	TL
Baseline: Most Frequent	29.7	4.23	3	28.4	14.08	5
Baseline: Random	4.20	4.21	71	6.36	15.04	71
System	37.3	10.52	47	37.3	16.77	52

Chunk-based Sentiment Representation: Narratives have patterns in ups and downs of sentiments [108]. Reagan et al. [87] showed that the pattern of changes in sentiments is significant for consumer experiences that results in success of stories. To capture such changes, we experiment with chunk-based sentiments and emotions representation. We divide the plot synopses into equally sized n chunks based on the word tokens and extract the sentiment and emotion features for each chunk. Then we run five-fold cross validation on the training data to observe the effect of chunk-based sentiments and emotions representation. Results in Table 5.3 show that dividing synopses into multiple chunks and using sentiment and emotion features from each chunk improves the performance of tag prediction. Although we observe noticeable improvements up to three chunks, TL remains similar where micro-F1 scores start to drop when we use more than three chunks. We suspect that higher number of chunks create sparseness in the representation of sentiments and emotions and hurts the performance. Therefore, we use sentiments and emotions features using three chunks in further experiments. As the chunk-based representation shows improvement in results, we plan to work capturing the flow of sentiments throughout the plots more efficiently in future work.

Table 5.3: Experimental results obtained by 5-fold cross-validation using chunk-based sentiment representations. Chunk-based sentiment features were combined with the other features described in Section 5.1.

Chunks	Top 3			Top 5		
	F1	TR	TL	F1	TR	TL
1	35.2	6.550	18.2	35.1	9.928	23.4
2	35.0	7.031	23.0	35.2	10.68	26.8
3	35.7	8.165	29.4	36.0	12.754	35.4
4	35.1	8.153	30.6	35.4	12.723	36.8
5	34.8	8.185	30.4	35.1	12.553	36.8
6	34.3	7.976	31.2	34.9	12.725	36.0

5.4 Conclusion

In this chapter, we described our first step towards high-level story understanding by the task of predicting a set of tags for movies using the written plot synopsis. We carefully designed a methodology that takes advantage of the lexical, semantics, and sentiments based characteristics of a narrative text to develop a One-versus-Rest machine learning model with logistic regression for the task. Experimenting with the hand-crafted linguistic features allowed us to gain significant insights about the problem of predicting tags for movies from the written plot synopses. We observed that the lexical features specifically character n-grams act as the most powerful ones among all the features. Comparatively weaker features can help to make the classifier better when combined with other features. Using sentiment and emotion-based features by diving the synopses into several chunks makes the model more effective, which gives us the motivation for more exploration into this direction.

Chapter 6

Modeling Flow of Emotions in Stories

Emotions are deeply connected to stories. A storyteller’s capability of playing with the emotions through the characters is one of the essential characteristics which makes a story unique, enjoyable, and close to the heart of the readers. Hence, we hypothesize that being able to model the emotional dynamics can help us to achieve a better understanding of stories, which in turn will benefit a tag prediction mechanism.

In this chapter, we design a neural method that simultaneously models the flow of emotions throughout a story and the textual representation from the written narrative to retrieve relevant tags. Throughout the rest of this chapter, we will first describe how we create an automatic representation of a narrative text with neural networks. Then we will discuss the method we follow to model the temporal emotional states in stories in this network. We will analyze our experimental results on tag prediction and investigate how this proposed design benefits the task.

6.1 Methods

In this section, we describe our proposed neural network based model that allows us modelling the emotion flow throughout stories along with the high-level textual representation from written synopses. Figure 6.1 shows the proposed model architecture. The proposed neural architecture has three modules. The first module (Section 6.1.1) uses a convolutional neural network (CNN) to learn high-level textual representations from synopses. The second module (Section 6.1.2) models the flow of emotions via a bidirectional long short-term memory (Bi-LSTM) network. And the last module contains hidden dense layers that operate on the combined representations generated by the first and second modules to predict the most likely tags for movies.

6.1.1 CNN Based Document Representation

Recent successes in different text classification problems motivated us to extract important word level features using convolutional neural networks (CNNs) [28, 52, 117, 43, 94]. We design a model that takes word sequences as input, where each word is represented by a 300-dimensional word embedding vector. We use randomly initialized word embeddings but also experiment with the FastText¹ word embeddings trained on Wikipedia using subword information. We decided to use this specific variant of word embedding models because of its ability of generating embeddings for out of vocabulary words as FastText computes a word embedding vector by using the vectors of the character n -grams in the particular word. We stack 4 sets of one-dimensional convolution modules with 1024 filters each for filter sizes 2, 3, 4, and 5 to extract word-level n -gram features [52, 117]. Each filter of size c is applied from window t to window $t + c - 1$ on a word sequence x_1, x_2, \dots, x_n . Convolution units of filter size c calculate a convolution output using a weight map W_c , bias b_c , and the ReLU

¹<https://fasttext.cc/docs/en/english-vectors.html>

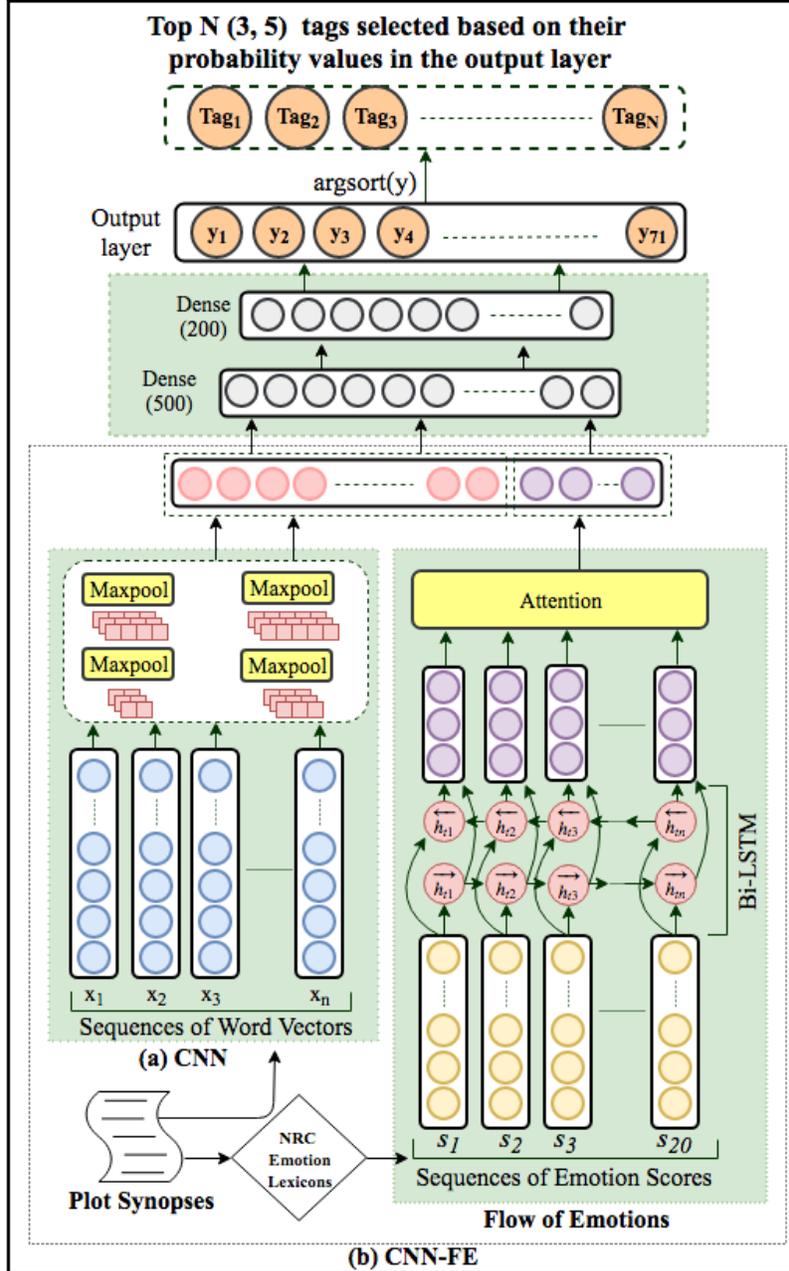


Figure 6.1: Convolutional Neural Network with Emotion Flow. The entire model is a combination of three modules. Module (a) learns feature representations from synopses using convolutional neural network. Module (b) incorporates emotion flows with module (a) to generate a combined representation of synopses. Module (c) uses these representations to predict the likelihood of each tag.

activation function [76]. The output of this operation is defined by:

$$h_{c,t} = ReLU(W_c x_{t:t+c-1} + b_c) \quad (6.1)$$

The ReLU activation function is defined by:

$$ReLU(x) = \max(0, x) \quad (6.2)$$

Finally, each convolution unit produces a high-level feature map h_c .

$$h_c = [h_{c,1}, h_{c,2}, \dots, h_{c,T-c+1},] \quad (6.3)$$

On those feature maps, we apply max-over-time pooling operation and take the maximum value as the feature produced a particular filter. We concatenate the outputs of the pooling operation for four filter sets that represent the feature representations for each plot synopsis.

6.1.2 Encoding Temporal Emotions

As we discussed in Section 2.1.2, stories can be described in terms of emotional shapes [108]; it has been shown that the emotional arcs of stories are dominated by six different shapes [87]. We believe that capturing the emotional ups and downs throughout the plots can help better understand how the story unfolds. For example, high concentration of *joy*, *trust*, and *anticipation* at the end of a story can indicate a happy ending. On the other hand, higher presence of *fear*, *anger*, and *disgust* could be the notion of violence and action. Being able to capture such properties in a story will enable us to predict relevant tags more accurately.

So we design a neural network architecture that tries to learn representations of plots using the vector space model of words combined with the emotional ups and downs of plots.

Human emotion is a complex phenomenon to define computationally. The Hourglass of Emotions model [16] categorized human emotions into four affective dimensions (*attention, sensitivity, aptitude, and pleasantness*), which started from the study of human emotions by [80]. Each of these affective dimensions is represented by six different activation levels that make up to 24 distinct labels called ‘elementary emotions’ that represent the total emotional state of the human mind. NRC² emotion lexicons [75] is a list of 14,182 words³ and their binary associations with eight types of elementary emotions from the Hourglass of Emotions model (*anger, anticipation, joy, trust, disgust, sadness, surprise, and fear*) with polarity. These lexicons have been used effectively in tracking the emotions in literary texts [73] and predicting success of books [66].

To model the flow of emotions throughout the plots, we divide each synopsis into N equally-sized segments based on words. For each segment, we compute the percentage of words corresponding to each emotion and polarity type (positive and negative) using the NRC emotion lexicons. More precisely, for a synopsis $x \in X$, where X denotes the entire collection of plot synopses, we create N sequences of emotion vectors using the NRC emotion lexicons as shown below:

$$x \rightarrow s_{1:N} = [s_1, s_2, \dots, s_N] \quad (6.4)$$

where s_i is the emotion vector for segment i . We experiment with different values of N , and $N = 20$ works better on the validation data.

As recurrent neural networks are good at encoding sequential data, we feed the sequence of emotion vectors into a bidirectional LSTM [36] with 16 units as shown in Figure 6.1.

²National Research Council Canada

³Version 0.92

This bidirectional LSTM layer tries to summarize the contextual flow of emotions from both directions of the plots. The forward LSTMs read the sequence from s_1 to s_N , while the backward LSTMs read the sequence in reverse from s_N to s_1 . These operations will compute the forward hidden states $(\vec{h}_1, \dots, \vec{h}_N)$ and backward hidden states $(\overleftarrow{h}_1, \dots, \overleftarrow{h}_N)$. For input sequence s , the hidden states h_t are computed using the following intermediate calculations:

$$\begin{aligned}
i_t &= \sigma(W_{si}s_t + W_{hi}h_{t-1} + W_{ci}c_{t-1} + b_i) \\
f_t &= \sigma(W_{sf}s_t + W_{hf}h_{t-1} + W_{cf}c_{t-1} + b_f) \\
c_t &= f_t c_{t-1} + i_t \tanh(W_{sc}s_t + W_{hc}h_{t-1} + b_c) \\
o_t &= \sigma(W_{so}s_t + W_{ho}h_{t-1} + b_o) \\
h_t &= o_t \tanh(c_t)
\end{aligned}$$

where W and b denote the weight matrices and bias, respectively. σ is the sigmoid activation function, and i , f , o , and c are *input gate*, *forget gate*, *output gate*, and *cell* activation vectors, respectively. The annotation for each segment s_i is obtained by concatenating its forward hidden states \vec{h}_i and backward hidden states \overleftarrow{h}_i , i.e. $h_i = [\vec{h}_i; \overleftarrow{h}_i]$. We then apply attention mechanism on this representation to get a unified representation of the emotion flow.

Attention models have been used effectively in many problems related to computer vision [72, 7] and have been successfully adopted in problems related to natural language processing [8, 92]. An attention layer applied on top of a feature map h_i computes the weighted sum r as follows:

$$r = \sum_i \alpha_i h_i \tag{6.5}$$

and the weight α_i is defined as

$$\alpha_i = \frac{\exp(\text{score}(h_i))}{\sum_{i'} \exp(\text{score}(h_{i'}))}, \quad (6.6)$$

where $\text{score}(\cdot)$ is computed as follows:

$$\text{score}(h_i) = v^T \tanh(W_a h_i + b_a) \quad (6.7)$$

where W , b , v , and u are model parameters. Finally, we concatenate the representation of the emotion flow produced by the attention operation and the output vector with the vector representation generated from the CNN module.

The concatenated vector is then fed into two hidden dense layers with 500 and 200 neurons. To improve generalization of the model, we use dropout with a rate of 0.4 after each hidden layer. Finally, we add the output layer \hat{y} with 71 neurons to compute predictions for 71 tags. To overcome the imbalance of the tags, we weight the posterior probabilities for each tag using different weight values. Weight value CW_t for tag $t \in T$ is defined by,

$$CW_t = \frac{|D|}{|T| \times M_t} \quad (6.8)$$

where, $|D|$ is the size of the training set, $|T|$ is the number of classes, and M_t is the number of movies having tag t in the training set. We normalize the output layer by applying a softmax function defined by,

$$\text{softmax}(\hat{y}) = \frac{\exp(\hat{y})}{\sum_{k=0}^{70} \exp(\hat{y}_k)} \quad (6.9)$$

Based on the ranking for each tag, we then select top N (3/5/10) tags for a movie.

6.2 Experiments

6.2.1 Data Processing and Training

As a preprocessing step, we lowercase the synopses, remove stop-words and also limit the vocabulary to top 5K words to reduce noise and data sparsity. Then we convert each synopsis into a sequence of 1500 integers where each integer represents the index of the corresponding word in the vocabulary. For the sequences longer than 1500 words, we truncate them from the left based on experiments on the development set. Shorter sequences are left padded with zeros.

During training, we use 20% of the training data as validation data. We tune various deep model parameters (dropouts, learning rate, weight initialization schemes, and batch size) using early stopping technique on the validation data. We use the Kullback-Leibler (KL) divergence [55] to compute the loss between the true and predicted tag distributions and train the network using the RMSprop optimization algorithm [99] with a learning rate of 0.0001. We implemented our neural network using the PyTorch deep learning framework.⁴

6.2.2 Baselines

We compare the model performance against three baselines: majority baseline, random baseline, and traditional machine learning system. The majority baseline method assigns the most frequent three or five or ten tags in the training set to all the movies. Similarly, the random baseline assigns randomly selected three or five or ten tags to each movie. Finally, we compare our results with the benchmark system created with the hand-crafted features. This benchmark system used different types of hand-crafted lexical, semantic, and sentiment

⁴<https://pytorch.org>

features to train a OneVsRest approach model with logistic regression as the base classifier.

6.2.3 Evaluation Measures

We try to follow the same evaluation methodology we followed for the hand-crafted feature based system. We create two sets of tags for each movie by choosing the most likely three and five tags by the system. Additionally, we report our results on a wider range of tags, where we select the top ten predictions. We evaluate the performance using the number of unique tags learned by the system (TL), micro averaged F1, and tag recall (TL). Tags learned (TL) computes how many unique tags are being predicted by the system for the test data (size of the tag space created by the model for test data). Tag recall represents the average recall per tag and it is defined by the following equation:

$$TR = \frac{\sum_{i=1}^{|T|} |R_i|}{|T|} \quad (6.10)$$

Here, $|T|$ is the total number of tags in the corpus, and R_i is the recall for the i^{th} tag.

6.3 Results

Table 6.1 shows our results for Top 3, Top 5, and Top 10 settings.

We will mainly discuss the results achieved by selecting top five tags as it allows us to compare with all the baseline systems and more tags to discuss about. As the most frequent baseline system assigns a fixed set of tags to all the movies, it fails to exhibit diversity in the created tag space. Still it manages to achieve a micro-F1 score around 28%. On the other hand, the random baseline system creates the most diverse tag space by using all of the

Table 6.1: Performance of tag prediction systems on the test data. We report results of two setups using three matrices (TL: Tags learned, F1: Micro f1, TR: Tag recall).

Methods	Top 3			Top 5			Top 10		
	TL	F1	TR	TL	F1	TR	TL	F1	TR
Baseline: Most Frequent	3	29.7	4.23	5	28.4	14.08	10	28.4	13.73
Baseline: Random	71	4.2	4.21	71	6.4	15.04	71	6.6	14.36
Baseline: Linguistic Features	47	37.3	10.52	52	37.3	16.77	—	—	—
CNN without class weights	24	36.8	7.99	26	36.7	12.62	27	31.3	24.52
CNN with class weights	49	34.9	9.85	55	35.7	14.94	67	30.8	26.86
CNN-FE	58	36.9	9.40	65	36.7	14.11	70	31.1	24.76
CNN-FE + FastText	53	37.3	10.00	59	36.8	15.47	63	30.6	26.45

possible tags. However its lower micro-F1 score of 6.30% makes it impractical to be used in real world scenario. At this point, we find an interesting trade-off between accuracy and diversity. It is expected that a good movie tagger will be able to capture the multi-dimensional attributes of the plots that allows to generalize a diverse tag space. Tagging a large collection of movies with a very small and fixed set of tags (e.g., majority baseline system) is not useful for either a recommendation system or users. Equally important is the relevance between the movies and the tags created for those movies. The hand-crafted features based approach [45] achieves a micro-F1 around 37%, which outperforms the majority and random baselines. But the system was able to learn only 52 tags, which makes 73% of the total tags. Our approach achieves a lower micro-F1 score than the traditional machine learning one, but it performs better in terms of learning more tags. We observe that the micro-F1 of the CNN model with only word sequences is very close (36.7%) to the hand-crafted features based system. However, it is able to learn only around 37% of the tags. By utilizing class weights in this model (see Eq. 8), we improve the learning for under-represented tags yielding an increase in *tag recall* (TR) and *tags learned* (TL). But the micro-f1 drops to 35.7%. With the addition of emotion flows to CNN, the CNN-FE model learns significantly more tags while micro-F1 and tag recall do not change much. Initializing the embedding layer

with pre-trained embeddings made a small improvement in micro-F1 but the model learns comparatively lesser tags. If we compare the CNN-FE model with the hand-crafted feature based system, micro-F1 using CNN-FE is slightly lower ($\approx 1\%$) than the feature based system. But it provides a strong improvement in terms of the number of tags it learns (TL). CNN-FE learns around 91% tags of the tagset compared to 73% with the feature based system. It is an interesting improvement, because model is learning more tags and it is better at assigning relevant tags to movies. We observe similar pattern for the rest of the two sets of tags where we select top three and ten tags. For all the sets, CNN-FE model learns the highest number of tags compared to the other models. In terms of micro-F1 and tag recall, it does not achieve the highest numbers but performs very closely.

6.4 Analysis

6.4.1 Incompleteness in Tag Spaces

One of the limitations of folksonomies is the incompleteness in tag spaces. The fact that users have not tagged an item with a specific label does not imply that that label does not apply to the item. Incompleteness makes learning challenging for computational models as the training and evaluation process penalizes the model for predicting a tag that is not present in the ground truth tags, even though in some cases it may be a suitable tag. For example, ground truth tags for the movie *Luther (2003)*⁵ are *murder*, *romantic*, and *violence*. And the predicted tags from our proposed model are *murder*, *melodrama*, *intrigue*, *historical fiction*, and *christian film*. The film is indeed a Christian film⁶ portraying the biography of Martin Luther, who led the Christian reformation during the 16th century. According to

⁵<http://www.imdb.com/title/tt0309820>

⁶<https://www.christianfilmdatabase.com/review/luther-2>

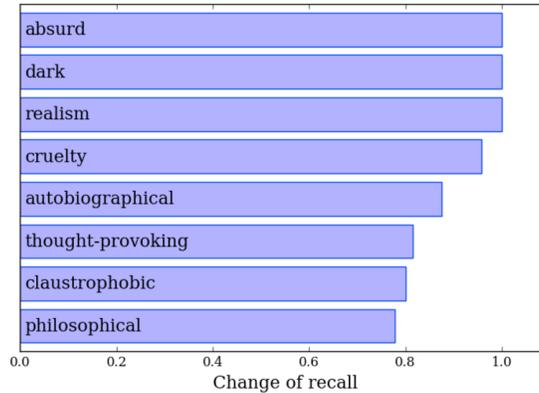


Figure 6.2: Tags with higher change of recall after adding the flow of emotions in CNN.

the Wikipedia, “*Luther is a 2003 American-German epic historical drama film loosely based on the life of Martin Luther*”.⁷ Similarly, Edtv⁸ (Table 6.2) has tags *romantic* and *satire* in the dataset. Our system predicted *adult comedy* and this tag is appropriate for this movie. In these two cases, the system will get lower micro-F1 since the relevant tags are not part of the ground truth. Perhaps a different evaluation scheme could be better suited for this task. We plan to work on this issue in our future work.

6.4.2 Significance of the Flow of Emotions

The results suggest that incorporating the flow of emotions helps to achieve better results by learning more tags. Figure 6.2 shows some tags with significant improvements in recall after incorporating the flow of emotions. We notice such improvements for around 30 tags. We argue that for these tags (e.g., *absurd*, *cruelty*, *thought-provoking*, *claustrophobic*) the changes in specific sentiments are adding new information helpful for identifying relevant tags. But we also notice negative changes in recall for around 10 tags, which are mostly related to the theme of the story (e.g., *blaxploitation*, *alternate history*, *historical fiction*,

⁷[https://en.wikipedia.org/wiki/Luther_\(2003_film\)](https://en.wikipedia.org/wiki/Luther_(2003_film))

⁸<http://www.imdb.com/title/tt0131369/>

sci-fi). It will be an interesting direction of future work to add a mechanism that can also learn to discern when emotion flow should contribute more to the prediction task.

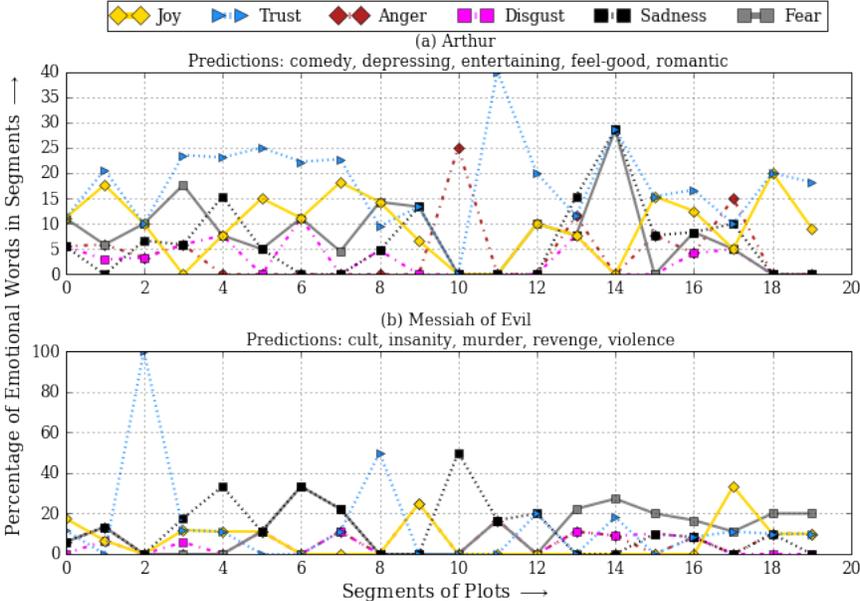


Figure 6.3: The flow of emotions in the plots of 2 different types of movies. Each synopsis was divided into 20 segments based on the words, and the percentage of the emotions for each segment was calculated using the NRC emotion lexicons. The y axis represents the percentage of emotions in each segment, whereas, the x axis represents the segments.

In Figure 6.3, we inspect how the flow of emotions looks like in different types of plots. Emotions like *joy* and *trust* are continuously dominant over *disgust* and *anger* in the plot of *Arthur (1981)*, which is a comedy film. We can observe sudden spikes in *sadness* and *fear* at segment 14, which is the possible reason for triggering the tag *depressing*. We observe a different pattern in the flow of emotions in *Messiah of Evil (1973)*, which is a horror film. Here the dominant emotions are *sadness* and *fear*. Such characteristics of emotions are helpful to determine the type and possible experiences from a movie. Our model seems to be able to leverage this information that is allowing it to learn more tags; specifically tags that are related to feelings.

6.4.3 Learning or Copying?

We found that only 11.8% of the 14,830 predicted tags for the $\sim 3\text{K}$ movies in the test data were found in the synopses themselves. 12.7% of the total 9,022 ground truth tags appear in the plot synopses. These numbers suggest that the model is not dependent on the occurrences of the tags in the synopses to make predictions, rather it seems it is trying to understand the plots and assign tags based on that. We also found that all the tags that were present in the synopses of the test data are also present in the synopses of the training data. Then we investigate what type of tags appear in the synopses and which ones do not. Tags present in the synopses are mostly genre or event related tags like *horror*, *violence*, *historical*. On the other hand, most of the tags that do not appear in the synopses are the tags that require a more sophisticated analysis of the plots synopses (e.g., *thought-provoking*, *feel-good*, *suspenseful*). It is not necessarily bad to predict tags that are in the synopses, since they are still useful for recommender systems. However, if this was the only ability of the proposed models, their value would be limited. Luckily this analysis, and the results presented earlier show that the model is able to infer relevant tags, even if they have not been observed in the synopses. This is a much more interesting finding.

6.4.4 Understanding Stories from Different Representations

Movie scripts represent the detailed story of a movie, whereas the plot synopses are summaries of the movie. The problem with movie scripts is that they are not as readily available as plot synopses. However, it is still interesting to evaluate our approach to predict tags from movie scripts. For this purpose, we collected movie scripts from our test set. We were able to find 80 movie scripts using the ScriptBase corpus [33].

We show the predicted tags from the synopses and scripts using our model in Table 6.2.

Table 6.2: Example of ground truth tags of movies from the test set and the generated tags for them using plot synopses and scripts.

<p>Title: A Nightmare on Elm Street 5: The Dream Child Ground Truths: cult, good versus evil, insanity, murder, sadist, violence Synopsis: cult, murder, paranormal, revenge, violence Script: murder, violence, flashback, cult, suspenseful</p>
<p>Title: EDtv Ground Truths: romantic, satire Synopsis: adult comedy, comedy, entertaining, prank, satire Script: comedy, satire, prank, entertaining, adult comedy</p>
<p>Title: Toy Story Ground Truths: clever, comedy, cult, cute, entertaining, fantasy, humor, violence Synopsis: comedy, cult, entertaining, humor, psychedelic Script: psychedelic, comedy, entertaining, cult, absurd</p>
<p>Title: Margot at the Wedding Ground Truths: romantic, storytelling, violence Synopsis: depressing, dramatic, melodrama, queer, romantic Script: psychological, murder, mystery, flashback, insanity</p>

In Table 6.3, we show the evaluation of tags generated using plot synopses and scripts. Despite having similar micro-f1 scores, *tag recall* and *tags learned* are lower when we use the scripts.

Table 6.3: Evaluation of predictions using plot synopses and scripts.

	Top 3			Top 5		
	F1	TR	TL	F1	TR	TL
Plot Synopses	29.3	8.04	28	38.7	15.70	35
Scripts	29.8	5.16	19	37.0	9.27	26

A possible explanation for this is the train/test mismatch since the model was trained using summarized versions of the movie, while the test data contained full movies scripts. Additional sources of error could come from the external info included in scripts (such as descriptions of actions from the characters or settings).

Table 6.4 shows that for most of the movies we generate very similar tags using the scripts and plot synopses. For 40% movies, at least 80% tags are the same. While the predictions

are not identical, these results show a consistency in the learned tags from our system. An interesting direction for future work would be to study what aspects in a full movie script are relevant to predict tags.

Table 6.4: Percentage of the match between the sets of top five tags generated from the scripts and plot synopses.

Percentage of Match	Percentage of Movies
$\geq 80\%$	40%
$\geq 40\% \ \& \ < 80\%$	47.5%
$\geq 20\% \ \& \ < 40\%$	11.25%

6.4.5 Challenging Tags

We found that these seven tags: *stupid*, *grindhouse film*, *blaxploitation*, *magical realism*, *brainwashing*, *plot twist*, and *allegory*, were not assigned to any movies in the test set. One reason might be that these are very infrequent (around 0.06% of movies have them assigned as their tags). This will obviously make them difficult to learn. Again, these are subjective as well. We believe that tagging a plot as stupid or brainwashing is complicated and depends on perspectives of a tagger. We plan to investigate such type of tags in the future.

6.5 Conclusion

In this chapter, we propose a method that learns word level feature representations from the written synopses using CNNs and models emotion flow throughout the stories using a bidirectional LSTM. We evaluated our model on the corpus containing plot synopses and tags of 14K movies. We compared our model against a majority and random baselines, and

the system built on traditional hand-crafted linguistic features. We found that incorporating emotion flows boosts prediction performance by improving the learning of tags related to feelings as well as increasing the overall number of tags learned.

Chapter 7

Hierarchical Representation of Narratives

In the previous two chapters, we demonstrated two different directions for designing mechanisms that take a narrative text (movie synopsis) as the input and generate a set of tags to characterize the underlying story at a high-level. In Chapter 5, we developed a machine learning model by extracting different traditional linguistic features from the texts. As the next step in Chapter 6, we have designed a sequential neural network for automatically learning high-level text representations and using that representation to generate tags. The flexibility of this network to capture the flow of emotions throughout the stories enhanced the network's capability to identify a higher number of story attributes than the feature-based approach.

In this chapter, we aim at improving such representation learning technique for stories and explore the direction of utilizing the hierarchical structure of documents rather than treating the documents as a plain sequence of tokens.

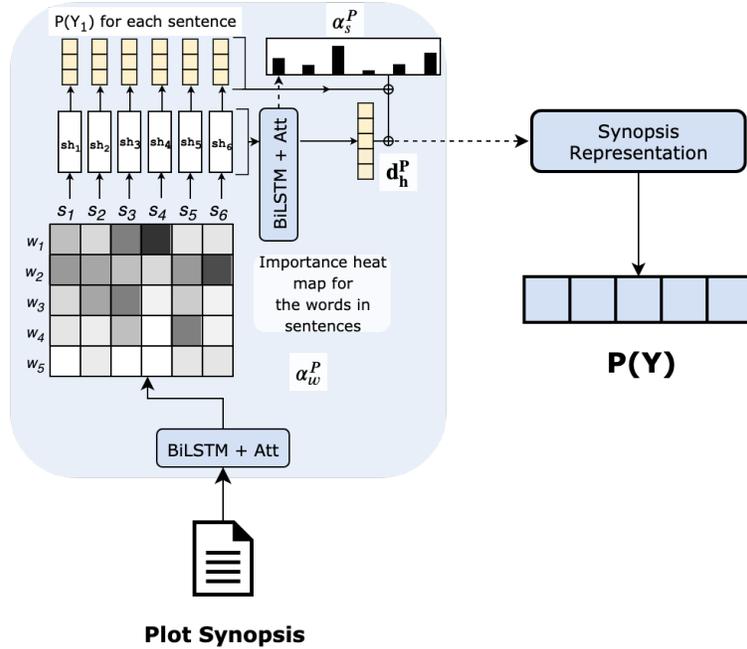


Figure 7.1: Architecture of the hierarchical model that learns weighted representation of words and sentences in a narrative text for creating a high-level story representation.

7.1 Model Architecture

Different words and sentences in a narrative text have different roles in the overall story. For example, some sentences narrate the setting or background of a story, whereas other sentences may describe different events and actions. Additionally, some sentences and words can be more helpful than others for identifying relevant tags from a written synopsis. With this consideration, we utilize an encoder similar to Yang et al. [116] that learns to weight important words and sentences and use this information as a step for creating a high-level document representation that enables better classification performance. Figure 7.1 shows the model diagram.

We represent a plot synopsis X_P consisting of L sentences (S_1, \dots, S_L) in a hierarchical structure instead of a long sequence of words like in Chapter 6. At first, for a sentence $S_i = (w_1, \dots, w_T)$ having T words, we create a matrix E_i where E_{it} is the vector representation

for word w_t in S_i . We use pre-trained word embeddings from Glove [78] to initialize E . Then, we encode the sentences using a bidirectional LSTM (BiLSTM) [36] with attention [8] mechanism. It helps the model to create a sentence representation \mathbf{sh}_i for the i_{th} sentence in X_P by learning to put more weight on the words that are crucial for making the correct predictions. The following equations show the transformations:

$$\begin{aligned}
\vec{h}_{w_{it}} &= \overrightarrow{LSTM}(\mathbf{E}_{it}), t \in [1, T] \\
\overleftarrow{h}_{w_{it}} &= \overleftarrow{LSTM}(\mathbf{E}_{it}), t \in [T, 1] \\
\mathbf{u}_{it} &= \tanh(\mathbf{W}_{wt} \cdot [\vec{h}_{w_{it}}, \overleftarrow{h}_{w_{it}}] + \mathbf{b}_w) \\
\mathbf{r}_{it} &= \mathbf{u}_{it}^\top \mathbf{v}_t \\
\alpha_{it} &= \frac{\exp(\mathbf{r}_t)}{\sum_t \exp(\mathbf{r}_t)} \\
\mathbf{sh}_i &= \sum_{t=1}^T \alpha_{it} \mathbf{h}_{it}
\end{aligned}$$

At the second step, we create a document representation from \mathbf{sh} , where the final prediction is made by jointly learning the tags from the document and individual sentence representations. We pass the encoded sentences \mathbf{sh} through another BiLSTM layer with attention to learn a document representation by weighting more weights on the important sentences. By taking the weighted sum of the hidden states and attention scores α_s^P for the sentences, we generate an intermediate document representation \mathbf{d}_h^P . Simultaneously, for each high level sentence representation \mathbf{sh}_i , we predict $P(Y)_{s_i}^P$, where Y represents the set of the predictable classes. This technique is inspired from Multiple Instance Learning (MIL), which we will describe in detail in Chapter 9. Then we weight $P(Y)_{s_i}^P$ by α_s^P and compute a weighted sum to prioritize the predictions made from comparatively important sentences. This sum is aggregated with \mathbf{d}_h^P to generate the final document representation. After generating this high level document representation of the synopses, We use a layer of

$|Y|$ linear units activated with the Softmax function to predict $P(Y)$, which represents the relevance of a tag y_i with the input text X_P .

7.2 Experiments

Dataset: We conduct our experiments on the MPST corpus proposed in Chapter 3, which has has around 15K movies and a set of tags. These tags were assigned to movies by the users of IMDB¹ and MovieLens.² The tagset contains close 70 tags that are representative of story related attributes (e.g., *thought-provoking*, *inspiring*, *violence*) and is free from any metadata (e.g., *cast*, *release year*). Additionally, synonymous tags (e.g., *suspense*, *suspenseful*, *full of suspense*, *tense*) are grouped into a single tag (e.g., *suspenseful*) to reduce noise in the tag space.

We tokenize the data using `spaCy`³ NLP library. To remove rare words and other noise, we retain the words that appear at least in ten synopses ($< .01\%$ of the dataset). Additionally, we replace the numeric tokens with a `cc` token. Through these steps we create a vocabulary of $\approx 42K$ word tokens. We represent the out of vocabulary words with a `<UNK>` token.

Implementation and Training We develop our experimental framework using PyTorch⁴ machine learning library and used the pre-trained model implementations by Hugging Face.⁵ We use KL divergence as the loss function for the network and Stochastic Gradient Descent (SGD) ($\eta = 0.2$, $\rho = 0.9$) to optimize the network parameters. We use a dropout rate of 50% between the layers and L2 regularization ($\lambda = 0.01$) to prevent overfitting. We observe faster convergence using batch normalization after each layer. For optimal performance, we tune parameters like the learning rate, regularization parameters. Based on $P(Y|X)$, we

¹<http://imdb.com>

²<http://movielens.org>

³spacy.io

⁴<http://pytorch.org>

⁵<https://huggingface.co/transformers/>

sort the tagset in descending order, so that the tags with higher probabilities are ranked at the top. Then in different settings we select the top N ($N=3, 5$) tags as the final tags to describe each movie.

Evaluation Metrics We aim at evaluating three aspects of the tag prediction models: a) accuracy of the top N predictions, b) diversity in the top N predictions, and c) performance in ranking most relevant tags at the top among the whole tag space. Regarding the first two aspects, we use micro-F1 and tags learned (TL) just as in Chapter 5 and Chapter 6. For the third aspect, we use *Multi-label Rank* (MLR) proposed in Chapter 4. This metric evaluates the overall prediction quality by considering the ranking of the target tags in models' predictions.

Baselines We compare our model against the following baselines:

1. **Top N:** This simple baseline ranks the tagset for a movie based on the frequency of tags in the training set. The most frequent $N(3, 5)$ tags are considered as the most relevant ones and assigned to movies.
2. **Hand-crafted linguistic features** Logistic regression based One-versus-Rest classifier trained on lexical, semantic, and sentiment features 5.
3. **Convolutional neural network with flow of emotion (CNN-FE)** Convolutional neural network based text encoder to extract features from written synopses and Bidirectional LSTMs to model the flow of emotions in the stories 6.
4. **Pre-trained language models** Large pre-trained language models (LM) built with Transformers [107] have shown impressive generalized performance in a wide range of natural language understanding (NLU) tasks like natural language inference, sentiment analysis, and question-answering in the GLUE benchmark [109]. We experiment with some of these language models like BERT [25] and RoBERTa [63] by treating them

as features with which we represent the synopses to predict tags. We also experiment with DistilBERT [89] that uses knowledge distillation [35] to reduce the model size of BERT while achieving similar performance in the NLU tasks.

We use the pre-trained tokenizers for creating the input representations for the pre-trained language models (WordPiece [114] for BERT and DistilBERT, and BPE [91] for RoBERTa). As the input length for these models is limited to 510 tokens, we experiment with a sliding window to divide the text into several chunks, extract LM features for each chunk, and combine them by a maxpooling operation. In other words, we create a window of size 510 and iteratively slide it over the tokenized texts to select multiple chunks having 510 tokens at most. For the simple sliding approach we use a stride of 510. As the chunks do not have any context in this approach, we experiment with an overlapping window of 20 tokens where the stride is 490. Then we wrap each of the chunks with special tokens ([CLS], [SEP]) If the last chunk has less than 512 tokens including the special tokens, we right pad it with 0s. We experiment by using LMs as a fixed feature extractor and also with fine-tuning the last Transformer layer. Due to the large memory requirement of these models for the long input texts, we are unable to experiment by fine-tune the whole model.

7.3 Results and Analysis

7.3.1 Quantitative Results

We present the detailed validation set results in Table 7.1 and summarized test results in Table 7.2. We mainly discuss the *Top 3* setting, where three tags are assigned to each instance by the systems. Among the language models, BERT (88.65) performed better than DistilBERT (86.00) and RoBERTa (86.59). In our experiments with sliding (S) and

Table 7.1: Results obtained on the validation set using different methodologies on the synopses. MLR and TL stand for *multi-label rank* and *tags learned* respectively. † indicates fine-tuning the last layer of BERT.

	MLR	Top3		Top5	
		F1	TL	F1	TL
Synopsis					
TopN	85.23	29.70	3	31.50	5
Features	85.23	36.90	40	36.80	48
CNN – FE	86.10	37.70	37	37.60	46
DistilBERT ^S	86.00	32.38	10	32.65	10
DistilBERT ^O	86.00	32.30	10	32.73	10
RoBERTa ^S	86.59	34.09	26	34.54	29
RoBERTa ^O	86.63	33.77	23	34.19	28
BERT ^S	88.65	35.84	28	36.33	31
BERT ^O	88.29	36.02	23	36.25	25
BERT ^{S†}	88.38	35.40	25	35.76	29
BERT ^{O†}	88.43	35.28	22	35.99	28
HN – Maxpool	89.36	36.72	17	36.19	28
HN – A	90.59	38.39	34	38.29	45
HN – A + MIL	91.32	38.54	49	38.99	54

overlapping (O) window techniques, sliding window performed better for BERT. We also try to fine-tune the last Transformer layer of BERT, with which we do not observe any improvement in performance. Hence, we proceed with pre-trained BERT with the sliding window approach. Our proposed hierarchical network with attention and Multiple Instance Learning (HN-A + MIL) achieved the best performance.

In the test set, multi-label Rank (MLR) achieved by BERT is 88.74, which is better than the other baseline methods. But F1 is 35.95 and the tags learned (TL) in the top three predictions for all movies are 26. We achieve an MLR of 89.32 and F1 of 36.31 when we use Maxpool in the hierarchical network (HN-Maxpool) to build the sentence and document representations. And, the number of distinct tags predicted from the tagset is 17, which is very low compared to the ground truth tagset size 71. The approach of utilizing word and sentence level importance via attention (HN-A) improves this number to 37

Table 7.2: Results obtained on the test set using different methodologies on the synopses. MLR and TL stand for *multi-label rank* and *tags learned* respectively.

	MLR	Top3		Top5	
		F1	TL	F1	TL
Synopsis					
Top N	85.66	29.70	3	28.40	5
Features	86.06	37.30	47	37.30	52
CNN – FE	86.51	36.90	58	36.70	65
BERT	88.74	35.95	26	36.41	32
HN – Maxpool	89.32	36.31	17	36.01	26
HN – A	90.45	37.90	37	37.67	46
HN – A + MIL	91.06	37.94	51	38.25	55

by enabling the model to better identify more fine-grained story attributes. It also shows improvement in terms of the other evaluation aspects (MLR=90.45, F1=37.90). Empowering the model with sentence-level predictions from the perspective of MIL shows further improvement (MLR=91.06, F1=37.94, TL=51). From these observations, we can conclude that creating word and sentence level representation hierarchically by learning to weight these representations can result in better performance for tag prediction for stories. Additionally, trying to predict the tags in the sentence level and aggregating these predictions to make the final prediction shows benefits in identifying various story attributes.

7.3.2 Inspecting the Hierarchical Representation

We analyze the reason behind the effectiveness of our proposed system by visualizing the attention weights at the sentence and word level for the synopsis of *Rush Hour 3* (see Figure 7.2). We can see that the sentences in the synopsis that describe important story events receive higher weights by the attention module. Similarly in the sentence level, important elements, events and characters (e.g., sword, Kenji, kills) are weighted more by the model. If we observe the tagsets provided in the caption and the highlighted words and sentences,

mystery	<u>30</u> Lee and Carter learn that Geneviève not only knows where the list is , she is the list .
action	<u>37</u> During a sword fight , Lee and Kenji fall off the tower and get caught in a safety net .
intrigue	<u>38</u> Kenji 's sword cuts the safety net open and it collapses , leaving both men hanging on for dear life .
dramatic	<u>41</u> Meanwhile , Carter single handedly defeats the rest of the Triad henchmen , unwittingly kills Jasmine , and saves Soo Yung .

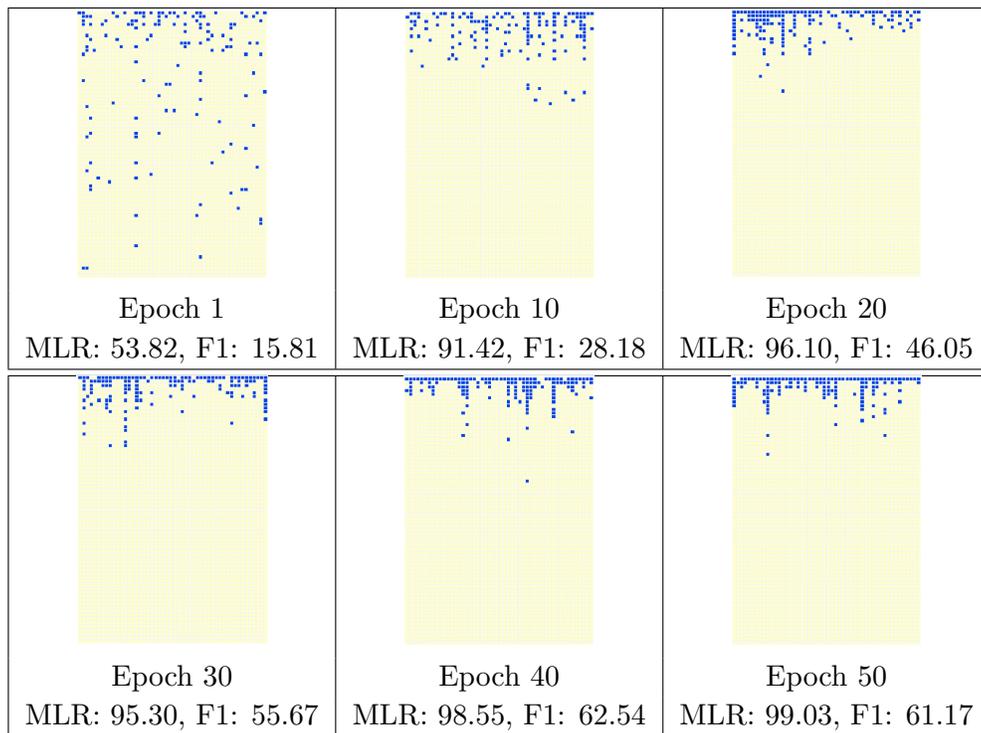
Figure 7.2: Example sentences from the synopsis of the movie **Rush Hour 3** with one of the most relevant tags from the sentence-level predictions. The importance of particular sentences and words for predicting tags are indicated by the highlight intensity of the sentence ids and words. Ground truth tags are *bleak, violence, comedy, murder*.

we can conclude that the model is efficiently modeling the correlations between the salient parts of the texts and the tags.

7.3.3 Inspecting Multi-label Rank

We present a visualization of the training progress with the update on Multi-Label Rank (MLR) and micro-F1 for top 3 predictions in Table 7.3 for a small subset of the data. The plots show that the target tags are progressively moved to the top with each new training epoch. But this training performance improvements are obscured with the f-measure, when most of the target tags are ranked at the top (after the 50th epoch), micro-F1 is only 61.17. But if we look at the MLR scores, they reflect the ranking performance more intuitively, which can eventually benefit any multi-label classification problems by providing a better way of performance evaluation.

Table 7.3: Visualization of the Multi-label Rank metric. Each cell presents a plot of the predictions after a particular epoch. In each plot, each column represents one data instance and each row represents the rank. There are $|Y|$ rows in total. The topmost cell in each column represents the class label having the highest probability score predicted by the model that makes it the top ranked class. Similarly the bottom-most cell is the lowest ranked class. If a cell represents a target class, the cell is colored as [Blue](#).



7.4 Conclusion

In this chapter, we proposed a model that exploits the weighted representation for individual words and sentences in a synopsis for constructing a high-level representation of the story. Such an approach showed better tag prediction performance than traditional feature based approach and sequential modeling of narratives. We also compared this system with pre-trained language models like BERT and the proposed model performed better than these language models. Analysis showed that, the effectiveness of this model comes from the ability of efficiently identifying important story elements, events, and characters. Finally we create tags from popular children stories, modern ghost stories, book summaries, and

TV shows using our model. These tags can accurately characterize the stories, which shows the generalizability and applicability of our model across different domains for generating high-level description of stories using tags.

Chapter 8

Applicability in Different Domains

As an application of high-level story characterization, so far we have explored the problem of describing movies with tags. We have developed a corpus that contains movie plot synopses with tags and used that corpus to design several methods for performing high-level story understanding to generate tags for movies. But such a mechanism could be applicable for other domains where storytelling is relevant. For example, automatically generated high-level tags can be employed to quickly describe fables, children stories, novels, TV series, or blogs. To verify the applicability, in this section, we use our model for generating tags for some widely popular children stories, books, and TV shows. We present these tags in Table 8.1.

We took some popular fairy-tales like the stories of *Cinderella*, *Rapunzel*, and *Aladdin* from the web and generated tags to describe them at a high-level. Although our model is trained on movie synopses to predict tags for movies, tags for these fairy-tales are surprisingly appropriate to describe the stories. For example, all of these fairy-tales are tagged as *fantasy*. Additionally, these tags often capture the *good versus evil*, *cute*, *whimsical*, and *romantic* characteristics of these stories.

At the next step, we experiment on a small set of *Modern Ghost Stories* collected from

Table 8.1: Generated tags from narratives that are not movie synopsis. We provide the full narrative texts in the Appendix C.

Children Stories	
Cinderella	fantasy cute romantic whimsical psychedelic
Snow White and the Seven Dwarfs	fantasy psychedelic romantic good versus evil whimsical
The Story of Rapunzel, A Brothers Grimm Fairy Tale	fantasy good versus evil psychedelic cute gothic
The Frog Prince: The Story of the Princess and the Frog	fantasy cute whimsical entertaining romantic
Aladdin and the Magic Lamp from The Arabian Nights	fantasy good versus evil action romantic whimsical
Modern Ghost Stories	
The Shadows on The Wall by Mary E. Wilkins Freeman	haunting gothic murder horror atmospheric
The Mass of Shadows By Anatole France	fantasy atmospheric gothic murder romantic
A Ghost By Guy De Maupassant	haunting flashback atmospheric murder paranormal
What Was It? By Fitz-James O'Brien	paranormal haunting gothic horror atmospheric
Story Summaries	
Romeo and Juliet by William Shakespeare	revenge murder romantic flashback tragedy
Harry Potter and Sorcerer's Stone by J. K. Rowling	fantasy good versus evil entertaining action comedy
Oliver Twist by Charles Dickens	murder revenge flashback romantic violence
The Hound of the Baskervilles by Arthur Conan Doyle	murder mystery gothic paranormal flashback
TV Series	
Game of Thrones Season 6 Episode 9: Battle of the Bastards	violence revenge murder action cult
The Big Bang Theory Season 3 Episode 22: The Staircase Implementation	romantic comedy flashback entertaining psychedelic
Friends Season 5 Episode 14: The One Where Everybody Finds Out	comedy entertaining adult comedy humor romantic
Narcos Season 1	murder neo noir violence action suspenseful

Project Gutenberg.¹ Table 8.1 shows that tags like *haunting*, *gothic*, *paranormal*, *murder*, and *horror* are generated for almost every story.

As processing large books is computationally expensive and our model is not trained to handle very large texts, we try to generate tags for some well-known novels from the synopses collected from the web. We notice that, most of the generated tags perfectly fit with the stories. For example, for *Harry Potter and Sorcerer's Stone*, some of the generated tags are *fantasy*, *good versus evil*, *entertaining*, which are undoubtedly appropriate to characterize this book. Ageless novel of detective Sherlock Homes, *The Hound of the Baskervilles* is also characterized by some well-fitted tags like *murder*, *mystery*, *flashback*, *paranormal*, and *paranormal*.

We observe the same trend when we describe some globally acclaimed TV shows like *Game of Thrones*, *The Big Bang Theory*, *Friends*, and *Narcos* with our model-generated tags. With the synopsis of what happened through the first season of *Narcos* (story of notorious drug kingpin Pablo Escobar), the model generated *murder*, *neo noir*, *violence*, *action*, and *suspenseful*, which accurately describe what the story offers. Additionally, highly rated episodes of *The Big Bang Theory* and *Friends* are tagged as *comedy*, *romantic*, *entertaining* and these tags are well-suited with these TV shows.

Observing these tags, we can conclude that despite being trained to generate tags for movies, our model is able to characterize the stories that are not movie plots. It indicates the promising generalization capability of our approach. Hence, our approach could be enhanced to be applied in other domains like books, children stories, TV shows to characterize the stories accurately at a high-level.

¹<http://www.gutenberg.org>

Part III

Employing Reviews for Story Characterization

Chapter 9

Retrieving Information from Reviews

In the second part of this dissertation, we focused on designing methods for high-level story characterization from narrative texts. More specifically, we experimented on different models to predict tags for movies to characterize the story at a high-level. In this second part, we will explore how can we take advantage of user reviews to enhance this tagset generation process. In this chapter, we will investigate two research questions:

- a) Can user reviews benefit a tag prediction system that generates tags for movie from the synopses?
- b) Is it possible to use such a prediction model to extract story descriptor tags from reviews in an unsupervised fashion?

We will start by discussing the motivation behind this work in detail. Then we will talk about the distinction between traditional opinion mining in product reviews and our approach for story descriptor extraction. We will describe the approach we design to investigate our research questions, and discuss the evaluation results.

Plot Synopsis

In late summer 1945, guests are gathered for the wedding reception of Don Vito Corleone's daughter Connie (Talia Shire) and Carlo Rizzi (Gianni Russo). ... The film ends with Clemenza and new caporegimes Rocco Lampone and Al Neri arriving and paying their respects to Michael. Clemenza kisses Michael's hand and greets him as "Don Corleone." As Kay watches, the office door is closed."

Review

Even if the viewer does not like mafia type of movies, he or she will watch the entire film ... Its about family, loyalty, greed, relationships, and real life. This is a great mix, and the artistic style make the film memorable.

violence action murder atmospheric
revenge mafia family loyalty
greed relationship artistic

Figure 9.1: Example snippets from plot synopsis and review of **The Godfather** and tags that can be generated from these.

9.1 Motivation

In a preliminary study we found that user reviews discuss the movie story line often, and in many cases this information combined with the synopses can lead to improved tag prediction. For example, in Figure 9.1 a reviewer writes that *The Godfather* is about *family, relations, loyalty, greed, and mafia*, whereas the gold standard tags from the plot are *violence, murder, atmospheric, action, and revenge*. Supporting information like this can provide more detailed information about a movie. Therefore, integrating the reviews with synopses in a tag prediction model can provide significant clue to a tag prediction model and improve the prediction performance.

However, waiting for reviews to accumulate in order to generate tags is not practical as not all the movies have enough user reviews to use in a model. Therefore, what we propose is a system that predicts tags by jointly modeling synopsis and reviews if available, but can still predict tags from a predefined tagset using only the synopsis if no reviews are available. Moreover, mining reviews and automatically extracting relevant tags outside of a gold set of tags could potentially alleviate the incompleteness problem commonly found in user generated tag spaces [48].

Mining tags to characterize stories can be modeled as a supervised opinion mining problem, which normally requires considerably large amounts of annotated tags in the reviews.

To get rid of this annotation burden, in this work we formulate the problem from the perspective of *Multiple Instance Learning* (MIL; Keeler and Rumelhart [49]) and develop a model that does not require tag level supervision in the reviews. Instead, our model learns to spot the opinion-heavy words in the reviews through document level supervision that are eventually transformed into a second tagset complementary to the gold tagset.

9.2 Opinion Mining vs Story Descriptor Extraction

Opinion mining is an effective technique for automatically identifying and summarizing user feedback about important attributes of products. It is widely applied on different product reviews and proven to be helpful for both manufacturers to improve their products and users to make calculated decisions before any purchase.

However, there is a subtle distinction between the reviews of typical material products (e.g., *phone, TV, furniture*) and story-based items (e.g., *literature, film, blog*). In contrast to the usual aspect based opinions (e.g., *battery, resolution, color*), reviews of story-based items often contain *end users' feelings, important events of stories, or genre related information*, which are abstract in nature (e.g., *heart-warming, slasher, melodramatic*) and do not have a very specific target aspect. Extraction of such opinions about stories has been approached by previous work using reviews of movies [118, 58] and books [60]. Such attempts are broadly divided into two categories. The first category deals with spotting words or phrases (*excellent, fantastic, boring*) used by people to express how they felt about the story and the second category focuses on extracting important opinionated sentences from reviews and generating a summary. In our work, while the primary task is to retrieve relevant attributes from the pre-defined set of tags, we also push for a system that can spot words in reviews capable of characterizing the story of movies.

Table 9.1: Statistics of the dataset. ^S denotes synopses and ^R denotes review summaries.

	Train	Val	Test
Instances	9746	2437	3046
Tags per instance	3	3	3
Reviews per movie	72	74	72
^S Sentence per document	50	53	51
^S Words per sentence	21	21	21
^R Sentence per document	117	116	116
^R Words per sentence	27	27	27

9.3 Dataset

We extended the MPST corpus (Chapter 3) by collecting up to 100 most helpful reviews per movie from IMDB. Out of the 15K movies in MPST, no reviews were found for 285 movies. The collected reviews often narrate the plot summary and describe opinions about movies. We noticed that reviews can be very long and sometimes contain repetitive plot summaries and opinions. Even some reviews can be uninformative about the story type. Moreover, the number of reviews for movies varies a lot, creating a challenge for modeling them computationally. So we summarize the reviews for a movie into a single document using Textrank¹ [69]. We observed that summarized reviews are usually free of repetitive information and aggregate the salient fragments from the reviews that are heavy with user opinions. All the plot synopses, reviews; tags are in English and Table 9.1 presents some statistics of the dataset.

9.4 Modeling

We model the task from the perspective of *Multiple Instance Learning* (MIL) and assume that each synopsis and review is a bag of instances (i.e. sentences in our task), where labels

¹We used the implementation from the Gensim library and converted the reviews into Unicode before summarizing.

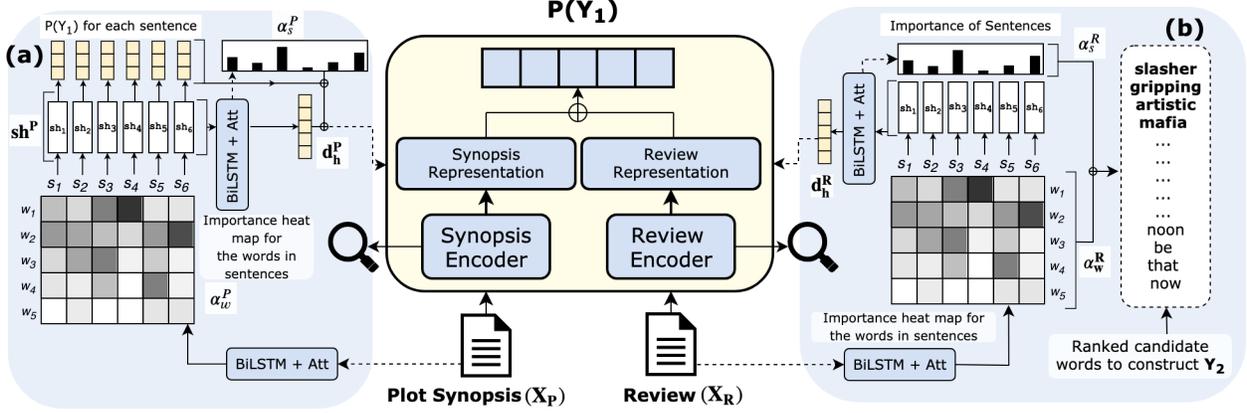


Figure 9.2: At the center, we show an overview of the model that takes a plot synopsis and a review summary as inputs, uses two separate encoders for the inputs to construct the high-level representations and uses them to compute $P(Y_1)$. **(a)** illustrates an enhanced view of the synopsis encoder. It uses a Bi-LSTM with attention to compute a representation \mathbf{sh}_i^P for the i th sentence in the synopsis. Additionally, a matrix of word-level attention weights α_w^P is generated that indicates the importance of each word in each sentence for correctly predicting $P(Y_1)$. Another attention based Bi-LSTM is used to create a synopsis representation d_h^P from the encoded sentences \mathbf{sh}_i^P . Additionally for each sh_i^P , sentence-level prediction $P(Y_1)_i$ is computed which is aggregated with d_h^P to create the final synopsis representation. **(b)** illustrates an almost similar encoder for reviews. To create an additional tagset Y_2 by mining opinions from the reviews, word-level importance scores α_w^R and sentence-level importance scores α_s^R are used). Apart from that, review representation d_h^R is computed in a similar way as in **(a)** which is used together with d_h^P to compute $P(Y_1)$.

are associated with each bag but not at the instance level. In such cases, a prediction is made for the bag by either learning to aggregate the instance level predictions [49, 27, 68] or jointly learning the labels for instances and the bag [40, 111, 53, 5, 115]. In our setting, we choose the latter. As we will show later, MIL improves prediction performance and also helps to investigate which parts of a story better correlate with certain tags.

Consider the input $X = \{X_P, X_R\}$ where X_P is a plot synopsis and X_R is a review summary. For a predefined target tagset $Y_1 = [y_1, y_2, \dots, y_{|Y_1|}]$, we would like to 1) model $P(Y_1|X)$ and 2) generate another tagset Y_2 complementary to Y_1 by extracting words from X_R . For this task, each sample is annotated with one or more tags from Y_1 but not with the tags in Y_2 . We use two separate encoder modules to generate high level representations \mathbf{d}_h^P and \mathbf{d}_h^R from X_P and X_R , and combine these representations to predict $P(Y_1)$. The overview

of our model is shown in Figure 9.2.

Inspired by our findings on hierarchical representation of narratives in the previous chapter, we employ two hierarchical document encoders for encoding synopses and reviews. In summary, we represent a plot synopsis X_P consisting of L sentences (S_1, \dots, S_L) in a hierarchical structure instead of a long sequence of words. At first, for a sentence $S_i = (w_1, \dots, w_T)$ having T words, we create a matrix E_i where E_{it} is the vector representation for word w_t in S_i . We encode the sentences using a bidirectional LSTM (BiLSTM) [36] with attention [8]. It helps the model to create a sentence representation \mathbf{sh}_i for the i_{th} sentence in X_P by learning to put more emphasis on the words that are crucial for making the correct predictions. At the second step, we create a document representation from \mathbf{sh} , where the final prediction is made by jointly learning the tags from the document and individual sentence representations. We pass the encoded sentences \mathbf{sh} through another BiLSTM layer with attention. By taking the weighted sum of the hidden states and attention scores α_s^P for the sentences, we generate an intermediate document representation \mathbf{d}_h^P . Simultaneously, for each high level sentence representation \mathbf{sh}_i , we predict $P(Y_1)_{s_i}^P$. Then we weight $P(Y_1)_{s_i}^P$ by α_s^P and compute a weighted sum to prioritize the predictions made from comparatively important sentences. This sum is aggregated with \mathbf{d}_h^P to generate the final document representation.

After generating the high level representation of the synopses and reviews we combine them together and use a layer of $|Y_1|$ linear units activated with the Softmax function to predict $P(Y_1)$. We experiment with two types of aggregation techniques: a) simple concatenation and b) gated fusion. In the first approach, we simply concatenate these two representations, whereas in the second approach we control the information flow from the synopses and review. While important story events and settings found from the synopses can correlate with some tags, viewers' reactions can also correlate with complementary tags. We believe that learning to control the contribution of information encoded from synopses

and reviews can improve the overall model performance. For instance, if the synopsis is not descriptive enough to retrieve relevant tags but the reviews have adequate information, we want the model to use more information from the reviews. Hence we use a gated fusion mechanism [6] on top of the encoded synopsis and review representations. For the encoded synopsis \mathbf{dh}_P and review representation \mathbf{dh}_R , the mechanism works as follows:

$$\begin{aligned} h_p &= \tanh(W_p \cdot dh_P) \\ h_r &= \tanh(W_r \cdot dh_R) \\ z &= \sigma(W_z \cdot [h_p, h_r]) \\ h &= z * h_p + (1 - z) * h_r \end{aligned}$$

9.5 Complementary Tag Generation from Reviews

Generating tags from reviews by extracting words can be seen from the perspective of MIL, where instance (i.e. word) level annotations are not present but each synopsis is annotated with some tags from the pre-defined tagset. We use this annotation for the synopses and force the model to learn to isolate the tags in the reviews by some weighting technique. For example in Figure 9.1, the blue tags are from the pre-defined tagset that was used to annotate the synopses and the yellow tags are induced from the reviews by our model without any annotation.

We observe that the model usually puts higher attention weights on opinion heavy words in the reviews. Therefore, we use the attention weights on words and sentences in reviews computed by the model to extract new tags not present in the tagset we use for training the models. Given $X = \{X_P, X_R\}$, we estimate $P(Y_1)$, that produces attention weight vectors α_W^R and α_S^R for X_R . For each word w_{ij} in sentence S_i in X_R , we compute an importance

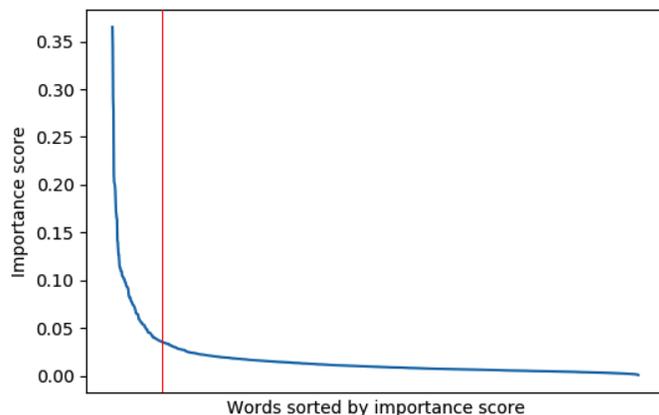


Figure 9.3: Illustration of the method for tag extraction from reviews using the importance score computed by the attention weights. All words left of the red vertical line are selected as candidate words.

score K_{ij} as:

$$K_{ij} = \alpha_{w_{ij}} \times \alpha_{S_i} \times |S_i| \quad (9.1)$$

Here, $\alpha_{w_{ij}}$ is the attention weight of word w_{ij} and α_{S_i} is the attention weight of the sentence. $|S_i|$ indicates the number of words in the sentence and it helps overcome the fact that word-level attention scores are higher in shorter sentences. We rank the words in the reviews based on their importance scores and choose the first few words as the primary candidate tags as shown in Figure 9.3. The idea is that the sorted scores create a downward slope and we stop at the point where the slope starts getting flat. We detect this by computing the derivative at this point based on its neighboring four points and set a threshold of 0.005 based on our observations on the validation set. After selecting the candidates we remove duplicates and words that are already in the predefined tagset to avoid redundancy. This method gives us a new tagset that is created from the opinions of users and provide new dimension to the story characterization task by extracting tags foreign to the tags in the dataset.

Table 9.2: Results obtained on the test set using different methodologies on the synopses and after adding reviews with the synopses. MLR and TL stand for *multi-label rank* and *tags learned* respectively. *: t-test with p -value < 0.01 .

	MLR	Top3		Top5	
		F1	TL	F1	TL
Synopsis					
Top N	85.66	29.70	3	28.40	5
Features	86.06	37.30	47	37.30	52
CNN – FE	86.51	36.90	58	36.70	65
BERT	88.74	35.95	26	36.41	32
HN(Maxpool)	89.32	36.31	17	36.01	26
HN(A)	90.45	37.90	37	37.67	46
HN(A) + MIL	91.06	37.94	51	38.25	55
Synopsis + Review					
Merge Texts	92.67	41.26	51	41.11	58
Concat Representations	92.68	40.64	55	40.82	62
Gated Fusion*	93.41	41.84	64	41.80	67
Subset (No Missing Reviews)					
Synopsis	91.05	38.24	50	37.99	55
Review	93.51	42.00	60	42.19	64
Both	93.52	42.11	65	42.00	68

9.6 Results

We report the results of our experiments on the test set in Table 9.2. We mainly discuss the *Top 3* setting, where three tags are assigned to each instance by the systems. Integrating reviews with the synopses seems to be effective as it improves the system performance in almost every aspect. As a simple baseline technique to combine reviews and synopses, we first add the review texts with the synopses to train a single hierarchical encoder based model. It shows improvements over the model that uses only the synopses (MLR=92.67, F1=41.26). Using two separate encoders for the synopses and reviews, concatenating the generated representations decreases MLR (92.68) and F1 (40.64), but increases TL (55). Combining these representations by gated fusion shows improvements and we achieve the best results so far (MLR=93.41, F1=41.84, TL=64). By performing a t-test, we found that

the gated fusion is significantly better (p -value < 0.01) than merging the texts and simple concatenation of the high-level representations of the synopses and reviews.

As we do not have reviews for ≈ 300 movies, we further experimented on a subset of the dataset by considering the movies having both synopsis and reviews. As shown in Table 9.2, reviews act as stronger data source than the synopses for classifying tags. Combining both does not affect much on the MLR (≈ 93.51) and F1 (≈ 42), but TL improves (60 vs 65). We found that combining synopses helps to identify tags like *plot twist*, *bleak*, *grindhouse film*, and *allegory*. It shows that our model is successfully capturing different story attributes from the reviews that are possibly difficult to find in synopses. Again, as reviews are not always available for movies, treating synopses as the primary data source and reviews as a complementary information is practical.

We analyze the reason behind the effectiveness of our proposed system by visualizing the attention weights at the sentence and word level for the reviews of *August Rush* in Figure 9.4. We can see that the sentences in the review that have user opinions characterizing the story receive higher weights. Similarly in the sentence level, important events and characters are weighted more by the model, and words in the review sentences that convey opinions about the storyline rather than the other aspects of the movie (e.g., music) receive more weights by the model.

9.6.1 Human Evaluation

In addition to quantitative evaluation, we also perform human evaluation to answer the following research questions: “*How relevant are our predicted tags compared to a baseline system? And how useful are the tags for the end users to get a quick idea about a movie?*”

We select CNN-FE 6 as the baseline system² to compare the quality of the tagsets created by

²We used the online demo system released by the authors to generate the tags.

5 The scenes have an appealing **fantasy** element , while at the same time , the plot manages to explore true - to - life human situations such as **bullying** of those who are different .

6 The **music is incredible** , and mostly consists of original scores .

7 It includes **gospel , rock and classical** , seamlessly integrated in a new way that works extremely well .

8 The plot is somewhat predictable and possibly a little " **sappy** " , but those elements are easily overcome by the moment - to - moment execution of the story .

Figure 9.4: Example sentences from the review of the movie **August Rush** with sentence ids and words highlighted based on importance on tag prediction. Ground truth tags are *thought-provoking, romantic, inspiring, flashback*.

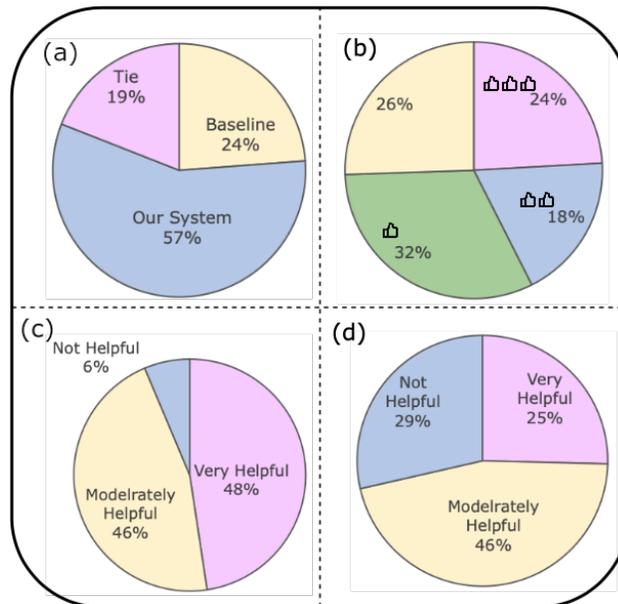


Figure 9.5: Summary of the human evaluation experiment results. (a) competing systems and the percentage of samples where their predictions were chosen as the best (b) complementary tags marked as helpful by number of annotators (c) helpfulness of the tagsets created from the closed set tagset (d) helpfulness of the complementary tagsets.

our system for 21 randomly sampled movies from the test set. For each movie, we instruct three human annotators to read the plot synopsis of a movie to understand the story. Then we show them two sets of tags for each movie and ask them to select the tags describing the story. In the first tagset, we mix the sets of five tags created by the baseline system and five tags from our system to avoid the possibility of bias towards any particular system. The second tagset is created by the complementary tags extracted from the user reviews (section 9.5). Then we ask annotators' opinion on whether these tagsets would help them decide if they want to see the movie or not. Finally, we reveal the title of the movie to them and ask if they have watched the movie or not and their opinion on the relevance of watching the film to fully evaluate the helpfulness of the tags. For this process, 21 people completed the evaluation and each of them evaluated three movies, for a total of 63 annotations.

We show the summary of the results in the pie charts in Figure 9.5.³ For each tag in the tagsets, we consider a tag as relevant if it is selected by at least two out of three annotators. With this criterion, our proposed system won⁴ for 12 out of 21 movies (57%), the baseline system won for five movies (24%), and the number of relevant tags were equal for the remaining four movies (19%). For these 21 movies, 141 tags (114 distinct tags) were extracted from the reviews in total. Out of these, 24% of the tags were selected as relevant by all of the three annotators, 18% tags by two annotators, and 32% tags by one annotator. It shows that around 74% of the extracted tags were marked as relevant by at least one annotator. When we ask the annotators about the helpfulness of the tagsets to decide whether to watch the movie or not, for the closed set tags, 6% annotations were mentioned as not helpful, 46% as moderately helpful, and 48% as very helpful. For the complementary tagsets created from the user reviews, 29% annotators mentioned them as not helpful, 46%

³Details in Appendix B.2.

⁴If annotators select more tags from System A than System B for a particular movie, we announce A as the winner.

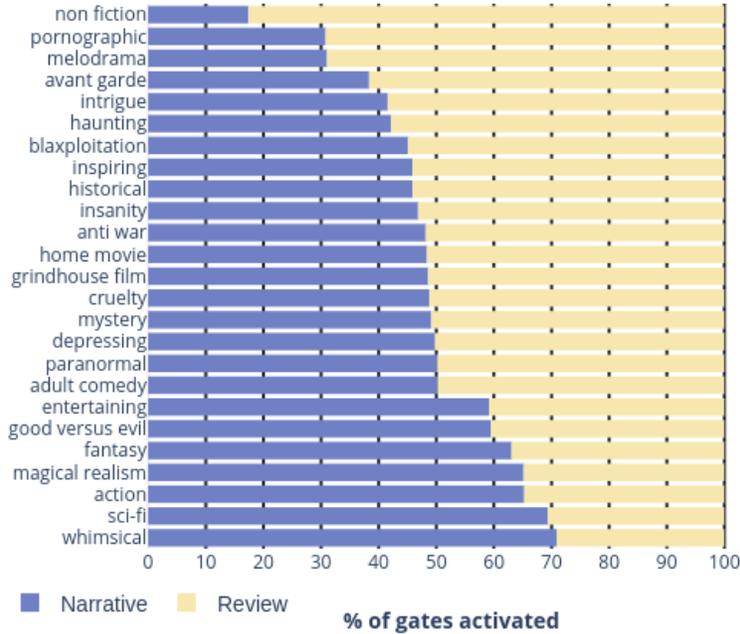


Figure 9.6: Percentage of gates activated for synopses and reviews. More active gates indicate more importance of the source for certain tags.

as moderately helpful, and 25% as very helpful. From these observations, we can conclude that our system is better at finding relevant story attributes than the baseline system, and in most of the cases the tags are helpful for the end users.

9.6.2 Information from Reviews and Synopses

By observing the predictions using only the synopses and having user reviews as an additional view, we try to understand the contribution of each view to identify story attributes. We notice that using user reviews improved ranking performance for tags like *non fiction*, *inspiring*, *haunting*, and *pornographic*. In Figure 9.6, we observe that the percentage of activated gates for the reviews was higher compared to synopses for the instances having the mentioned tags. Again, such tags are more likely to be related to the visual experience or somewhat challenging to understand from written synopses only. For example, synopses are more important to characterize *adult comedy* stories, but *pornographic* representation can

Table 9.3: Predicted tags from our model for new movies getting released in 2019 using the plot synopses. Bold-faced tags are already found to be assigned by users in IMDb.

Movie Title	Top 5 predictions
The Irishman	murder , neo noir, revenge , violence , flashback
Avengers: Endgame	good versus evil , fantasy, action , violence , flashback
Long Shot	entertaining, comedy , satire, humor, romantic
Annabelle Comes Home	paranormal, horror , gothic, cult, good versus evil
Once Upon a Time in Hollywood	comedy, violence , cult , humor , murder

be better identified by the viewers and this information can be easily conveyed through their opinion in reviews.

9.6.3 Predictions for New Movies

We collected plot synopses of some movies released in 2019 to generate some tags using our model and match them with the user assigned tags in IMDb. We present our predictions in Table 9.3 where we bold-face the tags that were also assigned by users. For example, our predictions for *The Irishman* are *murder*, *neo noir*, *revenge*, *violence*, *flashback*, where all the tags except *neo noir* were found in IMDb. Note that, accumulating reviews and tags is time consuming process and a large number of movies do not receive any reviews or tags at all. So we will check again in the coming months to see what tags appear for these movies. But this very small-scale experiment shows very promising results for generating relevant tags only from the synopses.

9.7 Conclusion

In this chapter we focused on characterizing stories by generating tags from plot synopses and reviews. As a first step to solve the problem, we extended an existing dataset of plot synopses and tags by collecting user reviews for movies. We model the problem from the

perspective of *Multiple Instance Learning* and developed a multi-view model architecture. The model learns to predict tags by identifying salient sentences and words of synopses and reviews. We empirically demonstrated that utilizing user reviews can further improve the performance and experimented on several methods for combining user reviews with synopses. Finally, we developed a methodology to extract user opinions that are helpful to identify complementary attributes of movies. We believe that this coarse story understanding approach can be extended to longer stories, i.e., entire books, and are currently exploring this path in our ongoing work.

Chapter 10

Conclusions

Being able to describe stories with a set of automatically generated tags can benefit many real life applications by letting people know what they can expect from any story-based items (e.g., movies, books) and making the selection process easier. However, computationally understanding the story from a narrative text, identifying the main concepts and attributes of the story, and generating a set of tags based on that high-level understanding is a challenging task in Natural Language Processing (NLP).

In this dissertation, we help to advance the field of NLP by contributing both resource and methodologies developed for automatic story characterization. Our first contribution is the *Movie Plot Synopses with Tags* (MPST) corpus, which is a large collection of movie plot synopses and their association with a set of tags that are capable of describing various attributes of the storyline of movies. This corpus is suitable for analysis of stories, information retrieval from narratives, and creative text generation. Additionally, it can serve as a good resource to support new research into multi-label classification in machine learning and NLP. We made this corpus publicly available to the research community and it has been actively used as a resource. As of March 2020, the dataset has been downloaded around 650 times.

As a next step, we presented several approaches to perform high-level story understanding from narrative texts which enable the generation of tags for describing stories. In our first proposed method, we showed how can we make use of traditional linguistic features extracted from written narratives to build a machine learning system for automatically generating tags. Through our analysis, we demonstrated that where lexical features are very powerful for accurate tag prediction, sentiments and emotions features are efficient for retrieving more diverse attributes of stories. Additionally, we showed that chunk-based representation is a more effective way for creating such feature representations. We made further development on this finding with a more robust method to model the emotional dynamics over time in stories. We showed that it is possible to model the flow of emotions with bi-directional recurrent neural networks (RNN) to create stronger story representations. Finally, we showed that exploiting the hierarchical structure of documents is necessary to create more powerful story representation. Our developed technique shows impressive generalized performance in tag generation for different types of stories (e.g., children books, TV series).

User reviews are a massive source of information and in this work, we showed that we can take advantage of such information for characterizing stories. Jointly modelling review summaries with stories can help a tag prediction model with supplementary information which improves tag prediction performance. Furthermore, such a technique allows us to extract new tags from user reviews in an unsupervised fashion.

To conclude, in this dissertation, we motivated the need for high-level story characterization systems that can improve the user experience while selecting story-based items like movies and books. The development process of such systems can enrich the field of Natural Language Processing by the innovation of techniques and tools for story analysis and interpretation. As part of this dissertation, we created new resources (dataset, source code, live

system) and made those resources publicly available to inspire more research and development in this area. We also envision several exciting scopes of further research. One of the possible future directions could be the invention of techniques for completely unsupervised extraction of story descriptor tags from reviews by utilizing the information found in stories. If we consider the story synopses as the cause and reviews as the effect, a possible system could be designed to identify the salient segments of the narratives and reviews and map the relation between them. This can help to spot the text fragments in the reviews that describe various story attributes. For example, events like kidnap, chase, fight, arrest in the synopses can be mapped to certain words in reviews like *suspenseful*, *action*, *mystery*. Being able to identify such relations can lead to a more exciting direction of fine-grained story analysis and mood-based story generation.

In this dissertation, we worked on story characterization systems that are specialized for generating tags for movies. Although our method can often generalize on stories from other domains, further efforts can enable this work efficiently applicable in numerous applications. Our work on high-level story understanding could be extended for enabling better analysis and understanding of long narrative texts and benefit the society in many possible ways. For example, such systems can be integrated into various recommendation systems to describe story-based items to users with tags and help the selection process easier. Future research could be conducted to explore if tags generated from stories can improve the performance of recommendation in movie and book recommendation engines. We believe, this work can be used in bibliotherapy to analyze why some books can treat particular psychological conditions. Such investigations can in turn help to develop text generation techniques for helping people to deal with different mental conditions. Finally, we are thrilled about the technologies (e.g., story analysis and understanding, story generation, tagging story-based

items) that can be developed with high-level story understanding and characterization process and we expect that the community will find our work useful on the way of these exciting innovations.

Bibliography

- [1] Apoorv Agarwal, Anup Kotalwar, and Owen Rambow. Automatic extraction of social networks from literary text: A case study on Alice in wonderland. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1202–1208, Nagoya, Japan, October 2013. Asian Federation of Natural Language Processing. URL <http://www.aclweb.org/anthology/I13-1171>.
- [2] Apoorv Agarwal, Anup Kotalwar, Jiehan Zheng, and Owen Rambow. Sinnet: Social interaction network extractor from text. In *The Companion Volume of the Proceedings of IJCNLP 2013: System Demonstrations*, pages 33–36, Nagoya, Japan, October 2013. Asian Federation of Natural Language Processing. URL <http://www.aclweb.org/anthology/I13-2009>.
- [3] Apoorv Agarwal, Sriramkumar Balasubramanian, Anup Kotalwar, Jiehan Zheng, and Owen Rambow. Frame semantic tree kernels for social network extraction from text. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 211–219, Gothenburg, Sweden, April 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/E14-1023>.
- [4] Gustavo Aguilar, Sudipta Kar, and Thamar Solorio. LinCE: A centralized benchmark for linguistic code-switching evaluation. In *Proceedings of the Twelfth International*

Conference on Language Resources and Evaluation (LREC 2020), Paris, France, May 2020. European Language Resources Association (ELRA).

- [5] Stefanos Angelidis and Mirella Lapata. Multiple instance learning networks for fine-grained sentiment analysis. *Transactions of the Association for Computational Linguistics*, 6:17–31, 2018. doi: 10.1162/tacl.a-00002. URL <https://www.aclweb.org/anthology/Q18-1002>.
- [6] John Arevalo, Thamar Solorio, Manuel Montes-y-Gómez, and Fabio A. González. Gated multimodal units for information fusion. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*, 2017. URL https://openreview.net/forum?id=S12_nqu0e.
- [7] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. *CoRR*, abs/1412.7755, 2014. URL <http://arxiv.org/abs/1412.7755>.
- [8] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2014. URL <http://arxiv.org/abs/1409.0473>.
- [9] Collin F Baker, Charles J Fillmore, and John B Lowe. The Berkeley Framenet project. In *Proceedings of the 17th international conference on Computational linguistics-Volume 1*, pages 86–90. Association for Computational Linguistics, 1998.
- [10] David Bamman, Brendan O’Connor, and Noah A. Smith. Learning latent personas of film characters. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 352–361, Sofia, Bulgaria,

August 2013. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P13-1035>.

- [11] Varshit Battu, Vishal Batchu, Rama Rohit Reddy Gangula, Mohana Murali Krishna Reddy Dakannagari, and Radhika Mamidi. Predicting the genre and rating of a movie based on its synopsis. In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation*, Hong Kong, 1–3 December 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/Y18-1007>.
- [12] Flavia Cristina Bernardini, Rodrigo Barbosa da Silva, Rodrigo Magalhaes Rodovalho, and Edwin Benito Mitacc Meza. Cardinality and density measures and their influence to multi-label learning methods. *Dens*, 1:1, 2014.
- [13] Douglas Biber. The multi-dimensional approach to linguistic analyses of genre variation: An overview of methodology and findings. *Computers and the Humanities*, 26(5):331–345, Dec 1992. ISSN 1572-8412. doi: 10.1007/BF00136979. URL <https://doi.org/10.1007/BF00136979>.
- [14] Kirk Borne. Collaborative annotation for scientific data discovery and reuse. *Bulletin of the American Society for Information Science and Technology*, 39(4):44–45, 2013.
- [15] Erik Cambria. An introduction to concept-level sentiment analysis. In *Mexican International Conference on Artificial Intelligence*, pages 478–483. Springer, 2013.
- [16] Erik Cambria, Andrew Livingstone, and Amir Hussain. The hourglass of emotions. In *Proceedings of the 2011 International Conference on Cognitive Behavioural Systems, COST’11*, pages 144–157, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-34583-8. doi: 10.1007/978-3-642-34584-5_11. URL http://dx.doi.org/10.1007/978-3-642-34584-5_11.

- [17] Erik Cambria, Soujanya Poria, Rajiv Bajpai, and Bjoern Schuller. Senticnet 4: A semantic resource for sentiment analysis based on conceptual primitives. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2666–2677, Osaka, Japan, December 2016. The COLING 2016 Organizing Committee. URL <http://aclweb.org/anthology/C16-1251>.
- [18] Gustavo Carneiro, Antoni B. Chan, Pedro J. Moreno, and Nuno Vasconcelos. Supervised learning of semantic classes for image annotation and retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(3):394–410, March 2007. ISSN 0162-8828. doi: 10.1109/TPAMI.2007.61. URL <http://dx.doi.org/10.1109/TPAMI.2007.61>.
- [19] Snigdha Chaturvedi, Shashank Srivastava, Hal Daum III, and Chris Dyer. Modeling evolving relationships between characters in literary novels. In Dale Schuurmans and Michael P. Wellman, editors, *AAAI*, pages 2704–2710. AAAI Press, 2016.
- [20] Snigdha Chaturvedi, Mohit Iyyer, and Hal Daumé III. Unsupervised learning of evolving relationships between literary characters. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [21] K. Choi, G. Fazekas, M. Sandler, and K. Cho. Convolutional recurrent neural networks for music classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2392–2396, March 2017. doi: 10.1109/ICASSP.2017.7952585.
- [22] Kenneth Ward Church and Patrick Hanks. Word association norms, mutual information, and lexicography. *Comput. Linguist.*, 16(1):22–29, March 1990. ISSN 0891-2017. URL <http://dl.acm.org/citation.cfm?id=89086.89095>.
- [23] Ido Dagan, Shaul Marcus, and Shaul Markovitch. Contextual word similarity and

- estimation from sparse data. In *Proceedings of the 31st Annual Meeting on Association for Computational Linguistics*, ACL '93, pages 164–171, Stroudsburg, PA, USA, 1993. Association for Computational Linguistics. doi: 10.3115/981574.981596. URL <https://doi.org/10.3115/981574.981596>.
- [24] Thierry Declerck, Nikolina Koleva, and Hans-Ulrich Krieger. Ontology-based incremental annotation of characters in folktales. In *Proceedings of the 6th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 30–34, Avignon, France, April 2012. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/W12-1006>.
- [25] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- [26] Sander Dieleman and Benjamin Schrauwen. Multiscale approaches to music audio feature learning. In *Proceedings of the 14th International Society for Music Information Retrieval Conference*, November 4-8 2013. http://www.ppgia.pucpr.br/ismir2013/wp-content/uploads/2013/09/69_Paper.pdf.
- [27] Thomas G. Dietterich, Richard H. Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artif. Intell.*, 89(1-2):31–71, January 1997. ISSN 0004-3702. doi: 10.1016/S0004-3702(96)00034-3. URL [http://dx.doi.org/10.1016/S0004-3702\(96\)00034-3](http://dx.doi.org/10.1016/S0004-3702(96)00034-3).

- [28] Cicero dos Santos and Maira Gatti. Deep convolutional neural networks for sentiment analysis of short texts. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 69–78, Dublin, Ireland, August 2014. Dublin City University and Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/C14-1008>.
- [29] Douglas Eck, Paul Lamere, Thierry Bertin-mahieux, and Stephen Green. Automatic generation of social tags for music recommendation. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 385–392. Curran Associates, Inc., 2008. URL <http://papers.nips.cc/paper/3370-automatic-generation-of-social-tags-for-music-recommendation.pdf>.
- [30] S. L. Feng, R. Manmatha, and V. Lavrenko. Multiple bernoulli relevance models for image and video annotation. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR'04*, pages 1002–1009, Washington, DC, USA, 2004. IEEE Computer Society. URL <http://dl.acm.org/citation.cfm?id=1896300.1896446>.
- [31] Marc Franco-Salvador, Sudipta Kar, Thamar Solorio, and Paolo Rosso. UH-PRHLT at SemEval-2016 Task 3: Combining lexical and semantic-based features for community question answering. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 814–821, San Diego, California, June 2016. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/S16-1126>.
- [32] Johannes Fürnkranz, Eyke Hüllermeier, Eneldo LozaMencía, and Klaus Brinker.

- Multilabel classification via calibrated label ranking. *Machine Learning*, 73(2):133–153, Nov 2008. ISSN 1573-0565. doi: 10.1007/s10994-008-5064-8. URL <https://doi.org/10.1007/s10994-008-5064-8>.
- [33] Philip John Gorinski and Mirella Lapata. Movie script summarization as graph-based scene extraction. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1066–1076, Denver, Colorado, May–June 2015. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/N15-1113>.
- [34] Amit Goyal, Ellen Riloff, and Hal Daumé, III. Automatically producing plot unit representations for narrative text. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP '10*, pages 77–86, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics. URL <http://dl.acm.org/citation.cfm?id=1870658.1870666>.
- [35] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. URL <http://arxiv.org/abs/1503.02531>.
- [36] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [37] Md. Zobaer Hossain, Md Ashrafur Rahman, Md Saiful Islam, and Sudipta Kar. A dataset for detecting fake news in Bangla. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, Paris, France, May 2020. European Language Resources Association (ELRA).

- [38] Andreas Hotho, Dominik Benz, Robert Jäschke, and Beate Krause. Ecml pkdd discovery challenge 2008 (rsdc'08). In *Workshop at 18th Europ. Conf. on Machine Learning (ECML'08)/11th Europ. Conf. on Principles and Practice of Knowledge Discovery in Databases (PKDD'08)*, volume 32, 2008.
- [39] Xiao Hu, J. Stephen Downie, and Andreas F. Ehmman. Lyric text mining in music mood classification. In *Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR 2009*, pages 411–416, 2009. ISBN 9780981353708.
- [40] Zhi hua Zhou, Yu yin Sun, and Yu feng Li. Multi-instance learning by treating instances as noni.i.d. samples. In *In Proceedings of the 26th International Conference on Machine Learning*, 2009.
- [41] Sheena Iyengar. *The art of choosing*. Twelve, 2010.
- [42] Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. Feuding families and former friends: Unsupervised learning for dynamic fictional relationships. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544. Association for Computational Linguistics, 2016. doi: 10.18653/v1/N16-1180. URL <http://aclanthology.coli.uni-saarland.de/pdf/N/N16/N16-1180.pdf>.
- [43] Sudipta Kar, Suraj Maharjan, and Tamar Solorio. RiTUAL-UH at semeval-2017 task 5: Sentiment analysis on financial data using neural networks. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 877–882, 2017.
- [44] Sudipta Kar, Suraj Maharjan, and Tamar Solorio. RiTUAL-UH at SemEval-2017

- Task 5: Sentiment analysis on financial data using neural networks. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 877–882, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/S17-2150. URL <https://www.aclweb.org/anthology/S17-2150>.
- [45] Sudipta Kar, Suraj Maharjan, A. Pastor López-Monroy, and Thamar Solorio. MPST: A corpus of movie plot synopses with tags. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France, May 2018. European Language Resources Association (ELRA). ISBN 978-2-9517408-9-1.
- [46] Sudipta Kar, Suraj Maharjan, and Thamar Solorio. Folksonomication: Predicting tags for movies from plot synopses using emotion flow encoded neural network. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2879–2891, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/C18-1244>.
- [47] Sudipta Kar, Gustavo Aguilar, and Thamar Solorio. Multi-view characterization of stories from narratives and reviews using multi-label ranking. *arXiv preprint*, 2019. URL <https://arxiv.org/pdf/1908.09083.pdf>.
- [48] Ioannis Katakis, Grigorios Tsoumakias, and Ioannis Vlahavas. Multilabel text classification for automated tag suggestion. In *Proceedings of the ECML/PKDD 2008 Discovery Challenge*, 2008.
- [49] Jim Keeler and David E. Rumelhart. A self-organizing integrated segmentation and recognition neural net. In J. E. Moody, S. J. Hanson, and R. P. Lippmann, editors, *Advances in Neural Information Processing Systems 4*, pages 496–503. Morgan-Kaufmann, 1992. URL <http://papers.nips.cc/paper/>

570-a-self-organizing-integrated-segmentation-and-recognition-neural-net.pdf.

- [50] Brett Kessler, Geoffrey Nunberg, and Hinrich Schutze. Automatic detection of text genre. In *8th Conference of the European Chapter of the Association for Computational Linguistics*, 1997. URL <http://aclweb.org/anthology/E97-1005>.
- [51] JungHyun Kim, Seungjae Lee, SungMin Kim, and Won Young Yoo. Music mood classification model based on arousal-valence values. In *Advanced Communication Technology (ICACT), 2011 13th International Conference on*, pages 292–295. IEEE, 2011.
- [52] Yoon Kim. 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, October 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D14-1181>.
- [53] Dimitrios Kotzias, Misha Denil, Nando de Freitas, and Padhraic Smyth. From group to individual labels using deep features. In *KDD*, 2015.
- [54] Vinodh Krishnan and Jacob Eisenstein. “You’re Mr. Lebowski, I’m the Dude”: Inducing address term formality in signed social networks. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1616–1626, Denver, Colorado, May–June 2015. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/N15-1185>.
- [55] S. Kullback and R. A. Leibler. On Information and Sufficiency. *Ann. Math. Statist.*,

- 22(1):79–86, 03 1951. doi: 10.1214/aoms/1177729694. URL <https://doi.org/10.1214/aoms/1177729694>.
- [56] Renaud Lambiotte and Marcel Ausloos. *Collaborative tagging as a tripartite network*, pages 1114–1117. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. ISBN 978-3-540-34384-4. doi: 10.1007/11758532_152. URL http://dx.doi.org/10.1007/11758532_152.
- [57] V. Lavrenko, R. Manmatha, and J. Jeon. A model for learning the semantics of pictures. In *Proceedings of the 16th International Conference on Neural Information Processing Systems, NIPS'03*, pages 553–560, Cambridge, MA, USA, 2003. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2981345.2981415>.
- [58] Fangtao Li, Chao Han, Minlie Huang, Xiaoyan Zhu, Ying-Ju Xia, Shu Zhang, and Hao Yu. Structure-aware review mining and summarization. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 653–661. Coling 2010 Organizing Committee, 2010. URL <http://aclweb.org/anthology/C10-1074>.
- [59] Xin Li, Lei Guo, and Yihong Eric Zhao. Tag-based social interest discovery. In *Proceedings of the 17th International Conference on World Wide Web, WWW '08*, pages 675–684, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-085-2. doi: 10.1145/1367497.1367589. URL <http://doi.acm.org/10.1145/1367497.1367589>.
- [60] E. Lin, S. Fang, and J. Wang. Mining online book reviews for sentimental clustering. In *2013 27th International Conference on Advanced Information Networking and Applications Workshops*, pages 179–184, March 2013. doi: 10.1109/WAINA.2013.172.

- [61] Marek Lipczak. Tag recommendation for folksonomies oriented towards individual users. In *In: Proc. of the ECML PKDD Discovery Challenge*, 2008.
- [62] Zachary Chase Lipton, David C. Kale, Charles Elkan, and Randall C. Wetzel. Learning to diagnose with LSTM recurrent neural networks. *CoRR*, abs/1511.03677, 2015. URL <http://arxiv.org/abs/1511.03677>.
- [63] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019. URL <http://arxiv.org/abs/1907.11692>.
- [64] Stuart Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [65] Suraj Maharjan, Sudipta Kar, Manuel Montes, Fabio A. González, and Thamar Solorio. Letting emotions flow: Success prediction by modeling the flow of emotions in books. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 259–265, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-2042. URL <https://www.aclweb.org/anthology/N18-2042>.
- [66] Suraj Maharjan, Sudipta Kar, Manuel Montes, Fabio A. Gonzalez, and Thamar Solorio. Letting emotions flow: Success prediction by modeling the flow of emotions in books. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 259–265, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/N18-2042>.

- [67] Nuno Mamede and Pedro Chaleira. Character identification in children stories. In *International Conference on Natural Language Processing (in Spain)*, pages 82–90. Springer, 2004.
- [68] Oded Maron and Aparna Lakshmi Ratan. Multiple-instance learning for natural scene classification. In *In The Fifteenth International Conference on Machine Learning*, pages 341–349. Morgan Kaufmann, 1998.
- [69] Rada Mihalcea and Paul Tarau. Textrank: Bringing order into text. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, 2004. URL <http://aclweb.org/anthology/W04-3252>.
- [70] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems*, NIPS’13, pages 3111–3119, USA, 2013. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=2999792.2999959>.
- [71] Gilad Mishne. Autotag: A collaborative approach to automated tag assignment for weblog posts. In *Proceedings of the 15th International Conference on World Wide Web*, WWW ’06, pages 953–954, New York, NY, USA, 2006. ACM. ISBN 1-59593-323-9. doi: 10.1145/1135777.1135961. URL <http://doi.acm.org.ezproxy.lib.uh.edu/10.1145/1135777.1135961>.
- [72] Volodymyr Mnih, Nicolas Heess, Alex Graves, and Koray Kavukcuoglu. Recurrent models of visual attention. *CoRR*, abs/1406.6247, 2014. URL <http://arxiv.org/abs/1406.6247>.
- [73] Saif Mohammad. From once upon a time to happily ever after: Tracking emotions in

- novels and fairy tales. In *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, LaTeCH '11, pages 105–114, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics. ISBN 9781937284046. URL <http://dl.acm.org/citation.cfm?id=2107636.2107650>.
- [74] Saif Mohammad and Peter D Turney. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34, Los Angeles, CA, June 2010. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/W10-0204>.
- [75] Saif M. Mohammad and Peter D. Turney. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465, 2013.
- [76] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pages 807–814, USA, 2010. Omnipress. ISBN 978-1-60558-907-7. URL <http://dl.acm.org/citation.cfm?id=3104322.3104425>.
- [77] Yoshiki Niwa and Yoshihiko Nitta. Co-occurrence vectors from corpora vs. distance vectors from dictionaries. In *Proceedings of the 15th Conference on Computational Linguistics - Volume 1*, COLING '94, pages 304–309, Stroudsburg, PA, USA, 1994. Association for Computational Linguistics. doi: 10.3115/991886.991938. URL <https://doi.org/10.3115/991886.991938>.
- [78] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.

- [79] Philipp Petrenz. Cross-lingual genre classification. In *Proceedings of the Student Research Workshop at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 11–21. Association for Computational Linguistics, 2012. URL <http://aclweb.org/anthology/E12-3002>.
- [80] Robert Plutchik. The Nature of Emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American Scientist*, 89(4):344–350, 2001.
- [81] V.I.A. Propp. *Morphology of the folktale: second edition*. American Folklore Society Bibliographical and Special Series. University of Texas Press, 1968. ISBN 9780292783768. URL <https://books.google.com/books?id=3Md3u9UPg0EC>.
- [82] Dheeraj Rajagopal, Erik Cambria, Daniel Olsher, and Kenneth Kwok. A graph-based approach to commonsense concept extraction and semantic similarity detection. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 565–570. ACM, 2013.
- [83] J. Read, B. Pfahringer, and G. Holmes. Multi-label classification using ensembles of pruned sets. In *2008 Eighth IEEE International Conference on Data Mining*, pages 995–1000, Dec 2008. doi: 10.1109/ICDM.2008.74.
- [84] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Mach. Learn.*, 85(3):333359, December 2011. ISSN 0885-6125. doi: 10.1007/s10994-011-5256-5. URL <https://doi.org/10.1007/s10994-011-5256-5>.
- [85] Jesse Read, Luca Martino, and David Luengo. Efficient monte carlo methods for multi-dimensional learning with classifier chains. *Pattern Recogn.*, 47(3):15351546,

- March 2014. ISSN 0031-3203. doi: 10.1016/j.patcog.2013.10.006. URL <https://doi.org/10.1016/j.patcog.2013.10.006>.
- [86] Jesse Read, Luca Martino, Pablo Olmos, and David Luengo. Scalable multi-output label prediction: From classifier chains to classifier trellises. *Pattern Recognition*, 48: 2096–2109, 06 2015. doi: 10.1016/j.patcog.2015.01.004.
- [87] Andrew J. Reagan, Lewis Mitchell, Dilan Kiley, Christopher M. Danforth, and Peter Sheridan Dodds. The emotional arcs of stories are dominated by six basic shapes. *CoRR*, abs/1606.07772, 2016. URL <http://arxiv.org/abs/1606.07772>.
- [88] Niloofar Safi Samghabadi, Deepthi Mave, Sudipta Kar, and Thamar Solorio. RiTUAL-UH at TRAC 2018 shared task: Aggression identification. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 12–18, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/W18-4402>.
- [89] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108, 2019.
- [90] Robert E Schapire and Yoram Singer. Boostexter: A boosting-based system for text categorization. *Machine Learning*, 39(2-3):135–168, 2000.
- [91] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1162. URL <https://www.aclweb.org/anthology/P16-1162>.

- [92] Paul Hongsuck Seo, Zhe Lin, Scott Cohen, Xiaohui Shen, and Bohyung Han. Hierarchical attention networks. *CoRR*, abs/1606.02393, 2016. URL <http://arxiv.org/abs/1606.02393>.
- [93] Mahsa Shafaei, Niloofar Safi Samghabadi, Sudipta Kar, and Thamar Solorio. Predicting the mpaa rating based on movie dialogues. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, Paris, France, May 2020. European Language Resources Association (ELRA).
- [94] Prasha Shrestha, Sebastian Sierra, Fabio Gonzalez, Manuel Montes, Paolo Rosso, and Thamar Solorio. Convolutional neural networks for authorship attribution of short texts. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 669–674, Valencia, Spain, 2017. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/E/E17/E17-2106.pdf>.
- [95] S.C. Sood, K.J. Hammond, S.H. Owsley, and L. Birnbaum. TagAssist: Automatic Tag Suggestion for Blog Posts. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM 2007)*, 2007.
- [96] Efstathios Stamatatos, Nikos Fakotakis, and George Kokkinakis. Automatic text categorization in terms of genre and author. *Computational Linguistics*, 26(4):471–495, 2000.
- [97] Martin Szomszor, Ciro Cattuto, Harith Alani, Kieron O’Hara, Andrea Baldassarri, Vittorio Loreto, and Vito D.P. Servedio. Folksonomies, the semantic web, and movie recommendation. In *4th European Semantic Web Conference, Bridging the Gap between Semantic Web and Web 2.0*, 2007. URL <http://eprints.soton.ac.uk/264007/1/ESWC2007.pdf>.

- [98] M. Tatu, M. Srikanth, and T. D’Silva. *RSDC’08: Tag Recommendations using Bookmark Content*. Workshop at 18th Europ. Conf. on Machine Learning (ECML’08) / 11th Europ. Conf. on Principles and Practice of Knowledge Discovery in Databases (PKDD’08), 2008. URL http://www.kde.cs.uni-kassel.de/ws/rsdc08/pdf/all_rsdc_v2.pdf.
- [99] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4(2):26–31, 2012.
- [100] Grigorios Tsoumakas and Ioannis Katakis. Multi-label classification: An overview. *Dept. of Informatics, Aristotle University of Thessaloniki, Greece*, 2006.
- [101] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Effective and efficient multilabel classification in domains with large number of labels. In *Proc. ECML/PKDD 2008 Workshop on Mining Multidimensional Data (MMD08)*, volume 21, pages 53–59. sn, 2008.
- [102] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Mining multi-label data. In *In Data Mining and Knowledge Discovery Handbook*, pages 667–685, 2010.
- [103] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Random k-labelsets for multilabel classification. *IEEE Transactions on Knowledge and Data Engineering*, 23(7):1079–1089, 2011.
- [104] Josep Valls-Vargas, Jichen Zhu, and Santiago Ontanón. Toward automatic role identification in unannotated folk tales. In *Proceedings of the Tenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 188–194. AAAI Press, 2014.

- [105] Menno van Zaanen and Pieter Kanters. Automatic mood classification using tf*idf based on lyrics. In *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR 2010, Utrecht, Netherlands, August 9-13, 2010*, pages 75–80, 2010. URL <http://ismir2010.ismir.net/proceedings/ismir2010-15.pdf>.
- [106] Thomas Vander Wal. Folksonomy definition, 2005. URL <http://www.vanderwal.net/folksonomy.html>.
- [107] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- [108] Kurt Vonnegut. *Palm Sunday: An Autobiographical Collage*. Random House Publishing Group, 2009. ISBN 9780307568069. URL https://books.google.com/books?id=Zd_9o3uyoVsC.
- [109] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of ICLR.*, 2019.
- [110] Changhu Wang, Shuicheng Yan, Lei Zhang, and Hong-Jiang Zhang. Multi-label sparse coding for automatic image annotation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1643–1650, June 2009.

- [111] X. Wei, J. Wu, and Z. Zhou. Scalable multi-instance learning. In *2014 IEEE International Conference on Data Mining*, pages 1037–1042, Dec 2014. doi: 10.1109/ICDM.2014.16.
- [112] Joseph Worsham and Jugal Kalita. Genre identification and the compositional effect of genre in literature. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1963–1973, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/C18-1167>.
- [113] Xi-Zhu Wu and Zhi-Hua Zhou. A unified view of multi-label performance measures. *CoRR*, abs/1609.00288, 2016. URL <http://arxiv.org/abs/1609.00288>.
- [114] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- [115] Yumo Xu and Mirella Lapata. Weakly supervised domain detection. *Transactions of the Association for Computational Linguistics*, 7:581–596, 2019. doi: 10.1162/tacl_a_00287. URL https://doi.org/10.1162/tacl_a_00287.
- [116] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1174. URL <https://www.aclweb.org/anthology/N16-1174>.

- [117] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, pages 649–657, 2015.
- [118] Li Zhuang, Feng Jing, and Xiao-Yan Zhu. Movie Review Mining and Summarization. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*, CIKM '06, pages 43–50, New York, NY, USA, 2006. ACM. ISBN 1-59593-433-2. doi: 10.1145/1183614.1183625. URL <http://doi.acm.org/10.1145/1183614.1183625>.

Appendices

Appendix A

Source Code of Multi-label Rank

```
1 def mlr(ground_truth_list, predicted_ranking):
2     """
3     The goal of this metric is to evaluate predictions in a
4     multi-label scenario by considering ranking, where,
5     1) Number of labels per sample is not fixed.
6     2) Labels for a single sample do not have any type of ordering
7     between themselves. That means change of order does not matter.
8     3) Considering that the ground truth labels stay at the top of
9     the ranking and create a block,
10    the idea is to check how far the labels are from the block
11    in the predictions.
12
13    Further illustration
14    =====
15
16    Let's say we have 10 labels in the dataset.
17    [A , B, C, D, E, F, G, H, I, J]
18
19    Ground Truth
20    A, D, G = G, D, A = D, G, A = D, A, G = A, G, D = G, A, D
21
22    Prediction 1 (Sorted by Probabilities)
23    [A , B, C, D, E, F, G, H, I, J]
24
25    >>> gt = [0, 3, 6]
26    >>> predicted_ranking = [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
27    >>> mlr(gt, predicted_ranking)
28    0.23809523809523805
```

```

29     >>> gt2 = [0, 2, 1]
30     >>> mlr(gt2, predicted_ranking)
31     0.0
32     >>> gt3 = [7, 8, 9]
33     >>> predicted_ranking = [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
34     >>> mlr(gt3, predicted_ranking)
35     0.8571428571428571
36     >>> gt4 = [7]
37     >>> predicted_ranking = [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
38     >>> mlr(gt4, predicted_ranking)
39     0.7777777777777778
40     >>> gt5 = [9]
41     >>> predicted_ranking = [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
42     >>> mlr(gt5, predicted_ranking)
43     1.0
44     >>> gt6 = [1]
45     >>> predicted_ranking = [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
46     >>> mlr(gt6, predicted_ranking)
47     0.1111111111111111
48     >>> gt7 = [1]
49     >>> predicted_ranking = [0, 1, 2]
50     >>> mlr(gt7, predicted_ranking)
51     0.5
52
53     """
54
55     block_len = len(ground_truth_list)
56     remaining_block_len = len(predicted_ranking) - block_len
57
58     distance = lambda pred_idx: max(0, pred_idx - block_len + 1) /
59                 remaining_block_len
60
61     score = sum(map(distance,
62                    [predicted_ranking.index(item)
63                     for item in ground_truth_list])
64                ) / block_len
65
66     return 1 - score
67
68
69 if __name__ == '__main__':
70     print(mlr('abc', 'abcdefg'))
71     print(mlr('abc', 'abdcgfe'))
72     print(mlr('abc', 'abdcefg'))
73     print(mlr('abc', 'abdefgc'))

```

```
74     print(mlr('abc', 'adefgbc'))
75     print(mlr('abc', 'begdafc'))
76     print(mlr('abc', 'defgabc'))
77
78     # For the examples in Chapter 4
79     print(mlr('abc', 'daefbcghijkl'))
80     print(mlr('abc', 'adefbghijckl'))
```

Appendix B

Human Evaluation

B.1 Questionnaire Interface

We employed 21 human subjects for the evaluation and each of them was asked to evaluate the generated tags for three movies. We developed a web interface using Flask web framework¹ to perform the human evaluation. In this section, we provide a walkthrough of the interface step by step using *Kronk's New Groove* as an example movie.

Before starting the evaluation task we verbally briefed them about the task, but also provide a written overview of the task and directions at the beginning of the web interface (Figure B.1). Then we presented the plot synopsis of a movie and asked them to read it (Figure B.2). At the next step, we presented the first tagset which is a mix of the tags produced by the baseline system in Chapter 6 and the model developed in Chapter 9. For this example movie, tagsets produced by these systems are (*cute, whimsical, feel-good, magical realism, prank*) and (*entertaining, romantic, flashback, prank, psychedelic*), respectfully. We provided the union of these tagsets to the user as shown in Figure B.3 and asked them to select the tags relevant to the story. Then we asked the subject if these tags are helpful

¹<https://palletsprojects.com/p/flask/>

Overview

In this task, you will be provided the plot synopsis of a movie and two sets of tags. If you read the plot synopsis, you will get an idea about the storyline of the movie. Then we will ask your opinion about the tagsets.

You have to look at the tags and think if the tags are related to the movie or helpful for you to decide if you want to watch the movie or not. For example, if you find the tag 'murder' in a tag set and there was an event of killing in the movie, **murder** is related to the movie. If you think that the movie has funny elements in the story, the tag **comedy** is relevant to the movie. Even if you do not want to watch a **comedy** movie, still the tag is relevant to the movie. It can be the case that plot synopsis doesn't explicitly mention the **murder** of someone, but you can feel that someone was murdered. In such cases, you can consider **murder** as a relevant tag. In a similar way, if you feel that there are **suspense** or **drama** in the movie, tags like **suspenseful** or **dramatic** are relevant tags for the movie.

Steps

- Read the plot synopses.
- Look at the tags from TAG SET 1.
- In the questionnaire select tags that are relevant. If all of the tags are relevant, you can select all. Do not select the options saying 'Ignore This!'.
- Look at the tags from TAG SET 2.
- Rate the helpfulness of TAG SET 2 in the scale of [Strongly Disagree, Somewhat Agree, and Strongly Agree]

Figure B.1: Questionnaire Interface Part 1: Written guideline to complete the task.

Movie 1

Please read the plot synopsis

Emperor Kuzco (David Spade) narrates the story about Kronk Pepikrankenitz (Patrick Warburton), now chef and Head Delivery Boy of Mudka's Meat Hut, is fretting over the upcoming visit of his father. Kronk's father always disapproved of young Kronk's culinary interests and wished that Kronk instead would settle down with a wife and a large house on a hill. In a flashback, Kronk tells the story of how he almost had both of these. As unwitting accomplice to Yzma (Eartha Kitt), the villainess of the first film who turned into a cat at the end of the original, but is now human again despite still having a tail, he goes along with her plan to sell sewer slime as a youth potion, he makes enough money to buy the old folks' home from the old folks and put his large new home there. Eventually Yzma is revealed as a fake and the old folks chase her down and corner her at a bridge over a river full of crocodiles. To prevent them from attacking her, she transforms herself into a rabbit, but is then caught and taken away by a condor. When Kronk realizes the old folks have sold everything they own in return for something which doesn't work, he gives his home back to them. Kronk, as camp counselor of the Junior Chipmunks at Camp Chippamunka he falls in love with fellow counselor Miss Birdwell (Tracey Ullman); but when one of his Chipmunks, Tipo, pulls a prank to win the camp championships and is caught, Kronk protects the boy at the cost of alienating his love. Kronk's father (John Mahoney) arrives and confusion ensues as several supportive friends try to pass themselves off to him as Kronk's wife and kids. But in the end Kronk realizes that his wealth is in his friendships, and this finally wins his father's thumbs up and Miss Birdwell's love. Meanwhile, just outside the house, Yzma is in the condor's nest with two eggs, which hatch and presumably attack her before the credits roll.

Figure B.2: Questionnaire Interface Part 2: Presenting a plot synopsis to the subject to read.

This is TAG SET 1. Select the tags from the following that are relevant to the story. (required)

1. entertaining
2. flashback
3. romantic
4. prank
5. psychedelic
6. cute
7. whimsical
8. feel-good
9. magical realism

TAG SET 1 is helpful to decide whether to watch the movie or not. (required)

Not Helpful
 Moderately Helpful
 Very Helpful

Figure B.3: Questionnaire Interface Part 3: Evaluating the relevance and helpfulness of the closed set tags.

This is TAG SET 2. Select the tags from the following that are relevant to the story. (required)

1. hilarious
2. kronk
3. moralising
4. funny
5. gags
6. cartoon

TAG SET 2 is helpful to decide whether to watch the movie or not. (required)

Not Helpful
 Moderately Helpful
 Very Helpful

Figure B.4: Questionnaire Interface Part 4: Evaluating the relevance and helpfulness of the open set tags extracted from the reviews.

Title of the movie is "Kronk's New Groove". Did you watch the movie? (required)

Yes

No

If you said YES, How much did watching the movie affect your judgement?

No extra benefit

Somewhat helped

It would be really hard without watching the movie

If you said NO, Do you think that watching the movie is necessary to know if the tags are relevant to the story?

Disagree

Somewhat Agree

Strongly Agree

Figure B.5: Questionnaire Interface Part 5: Inspection of any possible bias in the judgement.

How confident are you about your decision about TAG SET 1?(required)

Not confident

Moderate confident

Very confident

How confident are you about your decision about TAG SET 2?(required)

Not confident

Moderate confident

Very confident

Figure B.6: Questionnaire Interface Part 6: Collecting subjects confidence.

to decide whether to watch the movie or not in general. We repeated the similar process (Figure B.4) to evaluate the second tagset which was generated from the reviews by the extraction process described in Section 9.5. At the last step, we asked some questions to the subject to verify the reliability of the evaluation (See Figures B.5 and B.6).

B.2 Detailed Results

In this section, we present the detailed evaluation results of the generated tags by human subjects.

Table B.1: Data from the human evaluation experiment. **B** represents the tags predicted by the baseline system, **N** represents the tags predicted by our new system, and **R** represents the open set tags extracted from the user reviews by our system. If a tag is followed by a number in superscript, the number indicates the number of annotators who selected the tag as relevant to the story. We consider a tag as relevant if it has at least two votes. ♠ indicates the instances where our system’s predictions were more relevant compared to the baseline system, and ♣ indicates the opposite. For the rest of the instances, both systems had a tie. Annotators’ feedback about the helpfulness of the tagsets (closed set tags and open set tags) are presented by emoticons (😊: Very helpful, 😊: Moderately helpful, 😞: Not helpful). First three emoticons are the feedback from all the annotators for the tags from the baseline system and our system. Rest of the three emoticons are the feedback for the tags extracted from the user reviews.

IMDb id and Title	Tags with Number of Votes
tt0401398	B : feel-good ³ , magical realism ³ , cute ² , whimsical ¹ , prank ¹
Kronk’s New Groove	N : flashback ³ , romantic ³ , entertaining ² , prank ¹ , psychedelic
😊😊😊 [😞😊😊]	R : kronk ² , moralising ² , funny ² , cartoon ² , hilarious ¹ , gags ¹

tt0112887	B: pornographic ³ , adult comedy ¹ , neo noir ¹ , comic, blaxploitation
The Doom Generation	N: violence ³ , murder ³ , pornographic ³ , sadist ² , cult ¹
♠ [😬😬😬] [😬😬😬]	R: disturbing ³ , deranged ³ , tortured ³ , erotic ³ , violent ³ , psychotic ³ , goriest ³ , sexual ³ , kinky ² , weirdness ² , nihilistic ¹ , gay ¹ , homoerotic ¹ , irony, humourous, goth
tt0780606	B: plot twist ³ , murder ³ , suspenseful ² , intrigue ² , neo noir
Shock to the System	N: murder ³ , queer ³ , plot twist ³ , flashback ² , romantic ¹
[😬😬😬] [😬😬😬]	R: whodunit ³ , lesbian ³ , lgbt ³ , gay ³ , vengeance ³
tt0239949	B: comedy ³ , adult comedy ³ , humor ² , dramatic ² , entertaining ¹
Say It Isn't So	N: comedy ³ , romantic ² , humor ² , prank ¹ , entertaining ¹
♣ [😬😬😬] [😬😬😬]	R: humour ³ , funny ¹
tt0083869	B: neo noir ³ , adult comedy ² , humor ¹ , comedy ¹ , bleak ¹
Eating Raoul	N: murder ³ , adult comedy ² , pornographic ² , satire ¹ , comedy ¹
♠ [😬😬😬] [😬😬😬]	R: violent ³ , slapstick ² , humour ² , masochistic ¹ , bondage, kinky
tt0109650	B: dramatic ³ , historical ² , suspenseful ¹ , thought-provoking ¹ , neo noir ¹
Doomsday Gun	N: violence ³ , intrigue ³ , murder ² , flashback, alternate history
♠ [😬😬😬] [😬😬😬]	R: thriller ³ , cynical ² , backstabbing ¹ , conspiracy ¹ , amusing ¹ , evil ¹ , chases ¹ , paradox ¹ , nightmare ¹ , doomsday ¹ , chilling, mi6
tt0373283	B: historical ³ , action ³ , dramatic ³ , suspenseful ¹ , realism
Saints and Soldiers	N: violence ³ , historical ³ , murder ² , suspenseful ¹ , flashback
[😬😬😬] [😬😬😬]	R: massacre ³ , brutality ² , affirming ¹ , brotherhood, underbelly, christianity

tt0191423	B: entertaining ² , humor ¹ , comic ¹ , psychedelic, horror
Scooby-Doo	N: cult ¹ , flashback ¹ , comic ¹ , psychedelic, horror
Goes Hollywood	R: scooby ³
♣ [😞😞😞] [😞😞😞]	
tt0175059	B: good versus evil ³ , sci-fi ² , fantasy ² , alternate history ¹ , comic
Power Rangers Lost	N: good versus evil ³ , fantasy ² , violence ¹ , paranormal ¹ , psychedelic
Galaxy	R: mystical ² , mythic ¹ , cartoon, psycho, magical, funny
♣ [😞😞😞] [😞😞😞]	
tt0088805	B: thought-provoking ¹ , suspenseful ¹ , comic ¹ , paranormal, bleak
The Big Snit	N: psychedelic ³ , absurd ² , cult ¹ , philosophical ¹ , satire
♠ [😞😞😞] [😞😞😞]	R: surreal ³ , absurdist ² , existential ² , cartoon ¹ , demented
tt0064072	B: historical ³ , action ³ , dramatic ² , thought-provoking ¹ , anti war
Battle of Britain	N: historical ³ , flashback ² , violence ² , anti war, suspenseful
[😞😞😞] [😞😞😞]	R: gripping ³ , tragic ² , biographical ² , dogfights ¹ , sixties ¹
tt0067500	B: suspenseful ² , paranormal ² , murder ² , violence ² , revenge ¹
La noche del terror ciego	N: violence ² , murder ² , cruelty ² , cult ¹ , flashback ¹
♣ [😞😞😞] [😞😞😞]	R: disturbing ³ , satanic ¹ , gore ¹ , eroticism, lesbianism, visions, torture, tinged, subversive
tt0117913	B: revenge ³ , suspenseful ² , murder ² , violence ² , neo noir
A Time to Kill	N: revenge ³ , murder ² , violence ² , flashback ² , sadist ²
♠ [😞😞😞] [😞😞😞]	R: violent ³ , crime ³ , brutally ³ , vengeance ² , vigilante ² , sadism ¹ , poetic, depraved, fictional

tt0335345	B: dramatic ³ , thought-provoking ² , historical ² , suspenseful ¹ , allegory ¹
The Passion of the Christ	N: violence ³ , christian film ³ , murder ² , flashback ² , avant garde ¹
♠ [😊😊😊] [😊😞😊]	R: brutality ³ , symbolism ³ , slasher ² , treachery ² , enlightening ² , torture ² , lucid ¹ , occult ¹ , allusion ¹ , ironic
tt1185616	B: historical ² , thought-provoking ² , anti war ¹ , philosophical ¹ , alternate history
Vals Im Bashir	N: flashback ³ , violence ² , storytelling ² , murder ¹ , psychedelic
♠ [😊😊😊] [😊😞😞]	R: nightmares ³ , nightmare ³ , surreal ¹ , escapist ¹ , surrealism ¹ , disturbing, witty
tt2379386	B: action ³ , fantasy ³ , violence ³ , good versus evil ² , historical fiction ¹
In the Name of the King:	N: violence ³ , murder ³ , good versus evil ² , revenge ¹ , flashback ¹
The Last Mission	R: antihero ³ , magical ³ , campiness ¹ , dungeon ¹ , rampage ¹ , cinematic ¹ , masterpiece
♣ [😊😊😊] [😞😊😊]	
tt0085412	B: dramatic ³ , suicidal ¹ , realism ¹ , humor, thought-provoking
Deal of the Century	N: absurd ³ , comedy ² , satire ¹ , cult ¹ , humor
♠ [😞😊😊] [😞😊😊]	R: maniacal ³ , pathos ¹ , symbolism ¹
tt1355627	B: humor ² , clever ¹ , action ¹ , comic ¹ , thought-provoking
Evil Bong 2: King Bong	N: cult ² , comedy ² , violence ¹ , murder ¹ , revenge
♠ [😊😊😊] [😊😊😊]	R: humour ² , wicked ² , amusing ¹ , killer ¹ , evil ¹ , geeky ¹ , titular ¹ , laced, irreverence, homophobic
tt0023921	B: suspenseful ³ , murder ³ , revenge ² , sadist ¹ , neo noir
Cross Fire	N: murder ³ , violence ³ , suspenseful ³ , revenge ² , flashback
♠ [😊😊😊] [😞😊😞]	R: gunfight ³ , fistfights ¹ , classic

tt0154749	B: melodrama ² , romantic ² , flashback ² , intrigue ¹ , paranormal ¹
Kudrat	N: murder ² , flashback ² , romantic ² , revenge ² , paranormal ¹
♠ [😊😊😊] [😊😊😞]	R: thriller ³ , nightmares ² , reincarnation ² , chilling ¹ , karz, melancholy
tt0098575	B: romantic ³ , melodrama ² , historical fiction ¹ , queer, intrigue
Valmont	N: romantic ³ , revenge ² , murder ² , violence ² , flashback
♠ [😊😊😊] [😞😊😞]	R: cynicism ³ , irony ² , cruel ¹ , liaisons ¹ , humour ¹ , brutality, ruthless

In terms of the questions in Figure B.5, out of the 63 evaluations, in only one case the human subject had watched the movie and said that watching the movie earlier *somewhat* helped the judgement. In around 10% of the evaluations, the subjects *strongly* agreed that *watching the movie is necessary to know if the tags are relevant to the story*. 56% *somewhat* agreed with the statement and 34% *disagreed*. Hence, we assume that reading the plot synopses could let the subjects get an idea about the stories.

Appendix C

Out of Domain Stories

In this appendix, we provide the stories that we collected to generate high-level tags with our designed method in Chapter 8.

C.1 Children’s Story

C.1.1 Cinderella

Source: <https://www.thefablecottage.com/english/cinderella>

A long time ago there was a very beautiful girl named Cinderella.

Cinderella had long red hair, green eyes, and freckles all over her nose. She was clever and kind, and she loved to tell jokes.

But she was very unhappy. Her father and mother had died, and Cinderella lived with her stepmother and two stepsisters.

Although they all lived in a big house, they were actually quite poor. Their money was nearly gone.

Cinderella’s stepmother wanted one of her daughters to marry a rich man so they would no longer be poor.

But Cinderella’s stepsisters were not as pretty as Cinderella, not as kind as Cinderella, and not as funny as Cinderella.

The men who came to the house always fell immediately in love with Cinderella, and never noticed the stepsisters.

This frustrated the stepmother, so she ordered Cinderella to do all the chores.

“Sweep the hallway!” demanded the stepmother. “Clean the kitchen!” “Cook our dinner!” “Tidy our bedrooms!” “Mop the bathroom!” “Wipe the windows!” “Quick! Hurry!”

The stepmother tried hard to make Cinderella miserable.

The stepsisters had beautiful dresses and shoes, but Cinderella's dress was made of old rags.

The stepsisters always ate the most delicious foods, but Cinderella always ate scraps.

And the stepsisters slept in their comfortable beds in their bedrooms, but Cinderella slept on a straw bed on the kitchen floor.

The animals were Cinderella's only friends. In the evening she sat beside the fireplace in the kitchen and told jokes to the family of mice who lived in the wall. She talked to the cat.

"Things will get better one day," she told the cat. "Meow" the cat replied.

One day, while Cinderella was in the garden picking pumpkins, a letter arrived at the house. It was an invitation to the king's summer ball.

Cinderella's stepmother and stepsisters were very excited.

"The prince will be there!" "He is so handsome!" "He is so rich!" "He needs a wife!"

The stepsisters spent weeks preparing for the ball. They bought new dresses, new shoes, and new handbags.

On the day of the ball, Cinderella helped them put on their dresses and do their hair.

"Oh, I have a wonderful idea!" exclaimed the youngest stepsister. "Cinderella, come to the ball with us! It will be more fun if you are there!"

"Oh, but you don't have anything to wear" laughed the oldest stepsister. "You can't meet the prince wearing those dirty old clothes. What a shame. Maybe next time.

Cinderella tried not to cry. She finished dressing her sisters and then went down to the kitchen. She sat beside the fire and sighed.

"Things will get better one day," she told the cat. "Meow" the cat replied.

Just then, there was a flash of light, and an old woman appeared in a corner of the kitchen.

"Who who are you?" stammered Cinderella.

"I am your fairy godmother" said the old woman. "You are an orphan, and all orphans have a fairy godmother."

The fairy godmother stroked the cat.

"This cat tells me how kind you are. And how you always wish for things to get better one day. Today is that day, Cinderella. You are going to the king's ball. Fetch me a pumpkin!"

Cinderella ran into the garden and picked a big, orange pumpkin. The fairy godmother touched the pumpkin with her magic wand and it turned into a golden carriage.

"Come here, little mice!" she said to the mice in the wall. She waved her wand again, and the mice turned into six beautiful horses to pull the carriage.

"But I don't have a dress!" said Cinderella.

"Stand still," said the fairy godmother. She waved her wand again, and Cinderella's dirty clothes turned into a spectacular silver dress. Two beautiful glass shoes appeared on Cinderella's feet.

"Now go to the ball!" said the fairy godmother. "But you must be home by midnight! When the clock strikes twelve, your dress will turn back into rags, and your carriage will turn back into a pumpkin. ... Have fun!"

And with another flash of light, the fairy godmother disappeared.

"I'm going to the ball!" said Cinderella. "Meow" said the cat.

At the king's ball, the prince was very bored.

He felt like he had danced with every young lady in the kingdom. All of the ladies were wearing beautiful dresses, but none of the ladies were interesting. None of them understood his jokes.

The prince had just finished dancing with one of Cinderella's stepsisters when the room suddenly fell silent.

Everybody turned to look as the most beautiful girl walked through the door.

She had long red hair and kind green eyes. Her dress was silver. Her shoes sparkled, like they were made of glass.

It was Cinderella, but nobody recognised her. Not even her stepmother and stepsisters!

The prince's mouth dropped open. He had never seen a woman as beautiful as Cinderella.

He asked her to dance. They danced together all evening. The prince thought Cinderella was beautiful, but also kind, clever and funny. She laughed at his jokes, and he laughed at hers.

Cinderella was having such a wonderful time that she didn't notice that it was so late. The clock began to chime midnight.

Dong dong dong

"Oh no! I have to leave!" gasped Cinderella, and she ran out of the ballroom.

"Don't go! I don't even know your name!" shouted the prince. But Cinderella was already gone.

Cinderella fled from the palace so fast that she lost one of her glass shoes on the stairs.

When she got to the bottom of the stairs DONG! The clock finished chiming midnight.

Cinderella's beautiful dress turned back into rags, and her golden carriage turned back into a pumpkin.

"Drat," said Cinderella.

Just then, she saw the prince running towards her, holding the glass shoe she had dropped.

She did not want him to see her dressed in her dirty old rags. She felt ashamed, but there was nowhere to hide!

"Excuse me, miss," he said, panting. "Did you see where that beautiful girl went? This is her shoe! I must find her!"

The prince didn't recognise Cinderella without her beautiful clothes!

Cinderella shook her head. The prince ran off to continue his search. Cinderella walked all the way home.

Three weeks passed. The prince couldn't sleep. He could not stop thinking about the beautiful girl at the ball.

He waited for her to return to the palace, but she did not return. He waited for her to send a letter, but no letters came.

Finally, in desperation, he gave the glass shoe to a trusted messenger and ordered him to visit every house in the kingdom.

"Find the girl who this shoe belongs to, and bring her to me!"

The messenger visited hundreds of houses. At every house, women claimed the glass shoe was theirs. But when they tried on the shoe, their feet were too big, or too wide, or too small.

Finally, the messenger arrived at Cinderella's house. Cinderella's stepmother opened the door.

"Of course! Of course! Come in!"

She led the messenger into the dining room, where the two stepsisters were waiting.

The oldest sister said "Thank God! That's my shoe!" But when she tried the shoe, her foot was too wide.

The youngest sister said "Silly sister... It's not your shoe, it's my shoe!" But when she tried the shoe, her foot was too small.

The stepmother said "Get out of the way, girls, it's not your shoe. It's MY shoe!" and tried the shoe. But her foot was too long.

"Oh how silly!" said the stepmother. "The shoe must have shrunk in the rain..."

But the messenger was not so easily fooled. "Are there any other women in this house?" he asked.

"No one but our serving girl, and the shoe is certainly not hers..." said the stepmother.

"Fetch her immediately. Every woman in the kingdom must try the shoe," insisted the messenger.

When Cinderella arrived in the dining room she was wearing her usual rags, and her face was covered in dirt.

She put her dirty foot into the glass shoe and amazing! It wasn't too wide. It wasn't too long. It fit perfectly!

In a quiet voice she said, "It's my shoe."

"Please come with me," said the messenger. And before Cinderella's stepmother and stepsisters could stop them the messenger hurried Cinderella out the door and into a carriage.

Cinderella was taken to the palace to meet the prince. She was still wearing her dirty old dress, and her arms, legs and face were dirty. She looked at the floor because she felt so ashamed.

The prince took Cinderella's hand.

"Miss, please look at me," he requested kindly. And when she lifted her head and he saw her kind, green eyes, he knew that she was the girl he had fallen in love with at the ball.

They were married the next spring, and spent the rest of their lives laughing at each other's jokes.

C.1.2 Snow White and the Seven Dwarfs

Source: <https://www.storiestogrowby.org/story/snow-white-and-the-seven-dwarfs-bedtime-stories-for-kids/>

Once upon a time, a princess named Snow White lived in a castle with her father, the King, and her stepmother, the Queen. Her father had always said to his daughter that she must be fair to everyone at court. Said he, People come here to the castle when they have a problem. They need the ruler to make a fair decision. Nothing is more important than to be fair.

The Queen, Snow Whites stepmother, knew how much this meant to her husband. At the first chance, she went to her magic mirror. Mirror, mirror, on the wall, said the Queen. Who is the fairest of them all?

Snow White Story

Snow White is the fairest of them all! said the Magic Mirror.

What?! yelled the Queen. No one is more fair than I! The Queen must have the best of everything - everyone knows that. What could be more fair than that?

Snow White is the fairest of them all! repeated the Magic Mirror.

What do you know youre a mirror! roared the Queen. And she stormed off.

Still, the Queen was bothered. So bothered was she that the Queen decided to be rid of the girl, once and for all.

I cannot wait another day! she declared. The Queen called for her servant, a huntsman. Find a reason to take Snow White deep into the woods, she said, pointing her long finger at the servant. Then kill her.

Snow White Story

The huntsman was shocked! But she was the Queen and what could he do? The next day he took Snow White into the woods. As he drew his knife to slay her, Snow White turned around.

Look, she said, taking something out of her pocket. You have always been good to me. She held in front of him six perfect arrowheads that she had carefully shaped. Do you like them? she said. They are for you.

Snow White, said the huntsman. I cannot do this!

You can take these, said Snow White.

"That's not what I mean," said the servant. He dropped to his knees. How can I say this to you? The Queen, your step-mother, ordered me to kill you, he said. But I cannot!

She did what? Snow White called out with alarm.

You must run away! said the huntsman. Far into the woods. Now! And never come back to the castle!

Snow White turned and ran into the woods as fast as she could. Deeper and deeper she ran. It was getting dark, and the wolves were starting to howl. She tripped and her skirt was torn. Tall tree branches seemed to reach down to the very ground to grab her. She was scratched, bleeding and scared. Yet she ran on and on.

Then all of a sudden, far away, there was a light. Who was living so deep in the woods? She stepped up closer. It was a cottage! Yet no sound came from the cottage, only light from the windows.

Snow White and the Seven Dwarfs

Hello? she said, knocking softly on the door. Hello? No answer. The door was a little bit open. She opened it some more and stepped in. Hello, is anyone home?

She looked around. What a mess! She had never seen a messier living room.

This cottage may be the biggest mess I ever saw, she thought. But it's a roof over my head for tonight. Maybe if I clean up around here, I can earn my sleep.

As she cleaned, she thought of someone she already missed. Before her father had re-married, she and a Prince who lived in the next kingdom were getting to know each other. They would take long walks in the royal garden and tell each other stories, and laugh.

Snow White and the Seven Dwarfs

After the Queen had moved into the castle, her stepmother had made a new rule no more visitors. Now the Prince had to slip over the palace gate in secret. He would call out to her from under her window and they could talk a bit that way. It wasn't as good as the long walks but it was the best they could manage.

Now that she had to run away from home, would she ever see him again?

After Snow White cleaned up the living room, she went upstairs. On the second floor, there were seven little beds lined up in a row, as if for children. Tired from cleaning, Snow White yawned and lay across all seven of the beds. Soon she fell fast asleep.

In the meantime, the Seven Dwarfs were heading home from a long day of working in the jewel mines. When they opened the door, you can imagine their surprise when they saw their cottage all cleaned up!

Snow White and the Seven Dwarfs

What kind of magic is this? said one of the Dwarfs, whose name was Doc.

I wouldn't mind more magic like this! said another of the Dwarfs with a smile. His name was Dopey.

We'd better check upstairs, said another Dwarf, whose name was Grumpy. Something is fishy around here, that's for sure.

There lying across all their beds, was a young lady, fast asleep.

Who are you? said all the Dwarfs at once.

Snow White bolted awake. The Seven Dwarfs could tell she was as surprised as they were. Soon they all relaxed and shared their stories.

Snow White and the Seven Dwarfs

Snow White learned their names Bashful, Doc, Dopey, Grumpy, Happy, Sleepy, and Sneezy. She told them all about her step-mother. That her stepmother had tried to get the huntsman to kill her, that the huntsman had set her free in the woods, and that she could never go back home again.

Stay here, with us, said Bashful.

Thats sweet, said Snow White. But if I were to stay here at your home, I would have to do something for all of you.

You already cleaned up our place, said Sneezy.

Keeping the house clean will be easy, said Snow White, as long as we all pitch in. I will let everyone know what part they can do, and I will do my share too, of course.

Thats fair, said Happy.

But there must be something else I can do for you, said Snow White.

The Seven Dwarfs shrugged.

Do you know how to read? said Doc. "We have these books filled with wonderful tales and would love to be able to read them."

And so it was agreed that Snow White would give them reading lessons.

Snow White and the Seven Dwarfs

To celebrate their new friendship, Snow White and the Seven Dwarfs sang and danced the night away.

The next morning before they left for work, the Seven Dwarfs warned Snow White she must not open the door to anyone. After all, who knows what evil her stepmother might do? The princess nodded in agreement, and the Dwarfs left the house. The princess prepared her first reading lesson. She also prepared a good hot meal for the Seven Dwarfs when they returned home that night. And so the days passed.

Back at the castle, the Queen marched up to her mirror. Mirror, mirror on the wall," she demanded. "Who is the fairest of them all?

Snow White is the fairest of them all! said the Magic Mirror.

Snow White Story

Thats impossible! screamed the Queen. The girl is no longer alive!

Snow White lives! said the Magic Mirror. And an image was shown on the mirror of Snow White living in the cottage of the Seven Dwarfs.

The Queen turned red with rage. She screamed, She will not get away with this!

At the cottage of the Dwarfs the next afternoon, when the Seven Dwarfs were away at work, there was a knock on the door.

Who is it? said Snow White. She remembered the warning of the Seven Dwarfs not to open the door to anyone.

Its only a poor old woman, came a squeaky voice, selling apples. Yet it was the evil Queen, disguised as an old woman. Its raining out here, my dear, said her voice through the door. Please let me in.

Poor thing," thought Snow White, "having to go door to door selling apples in the rain." And so she opened the door.

Take a look at this big red apple, said the old woman, who as you know by now was really the Queen in disguise. She held the red apple close to Snow Whites face. Lovely, my dear, isnt it?

Snow White and the Seven Dwarfs

I would like very much to buy your apple, said Snow White. But I have no money.

That fine comb in your hair will make a good trade, said the old woman.

Well, all right then! said Snow White. She took the comb out of her hair and gave it to the old woman, who then gave her the apple. Snow White took a big bite. Alas, the fruit was poisoned! At once, Snow White fell to the ground in a deep sleep.

YES! shouted the Queen, pumping the air with her fists.

Just then the door flew open. In marched the Seven Dwarfs, home from the day's work. Shocked indeed they were to find Snow White lying on the floor and what must be her stepmother beside her, laughing!

They chased that evil Queen out the door, and into the storm. Up to the very top of a mountain they chased her. All of a sudden, lightning hit the mountain! The Queen fell, and she was never seen again.

But there was nothing to help poor Snow White. She stayed absolutely still in her deep sleep. The Seven Dwarfs gently lifted her into a glass coffin. Day and night they kept watch over her.

One day, the Prince happened to pass through. Ever since he had learned that Snow White was missing at the castle, he was searching for her, far and wide. Now he had finally found her, but in such a state! The Prince pulled open the glass coffin. Her face seemed so fresh, even in that deep sleep.

Snow White and the Seven Dwarfs

He gently took one of Snow White's hands in his own and kissed it. At once, Snow Whites eyes opened! With Love's First Kiss, the evil Queen's spell was forever gone. Now nothing stood in the way for Snow White and the Prince to be together forever. They returned to the kingdom and lived happily ever after.

C.1.3 The Story of Rapunzel, A Brothers Grimm Fairy Tale

Source: <https://www.storiestogrowby.org/story/early-reader-rapunzel-fairy-tale-story-kids/>

Chapter 1 The Carpenter and His Wife

ONCE UPON A TIME, there lived a carpenter and his wife. More than anything, they wanted a child of their own. At long last, their wish came true the wife was going to have a baby!

From the second floor window of their small house, the wife could see into the garden next door. Such fine fresh rows of plants and flowers there were! But no one dared to go over the garden wall to see them up close. For the garden belonged to a witch!

Rapunzel Story

One day the wife was looking down at the garden from her window. How fresh-looking were those big green heads of lettuce! It is just what I need to eat! said the wife to her husband. You must go and get me some.

But we cannot! said the carpenter. You know as well as I do that the garden belongs to the witch, who lives next door.

If I cannot have that lettuce, said the wife, I will not eat anything at all! I will die!

What could the carpenter do? Late that night, he climbed over the garden wall. With very quiet steps, he took one green head of lettuce. With more quiet steps, he went back over the garden wall. His wife ate up the lettuce right away.

But eating the lettuce only made her want more! If she could not have more lettuce, she said, there was nothing she would eat at all! So the next night, the carpenter climbed back over the garden wall. He picked up one more head of lettuce. All at once came a high, loud, voice.

STOP! What do you think you are doing?

I am getting lettuce for my wife, said the carpenter.

You thief! yelled the witch. You will pay for this!

Please! said the carpenter. My wife is going to have a baby. She saw your lettuce and wanted it so very much.

Why should I care about that? shouted the witch.

I will do anything! said the carpenter. He thought, Maybe I can build her something.

You say you will do anything? said the witch.

Yes, he said.

Fine! said the witch. Here's the deal. Go ahead take all the lettuce you want. Your wife will have a baby girl. And when she does, the baby will be mine!

What?! said the carpenter. I would never agree to that!

You already did! said the witch. And she laughed an evil laugh.

Chapter 2 The Tower

Soon the wife had a baby girl, just as the witch had said. To keep the baby safe from the witch, the carpenter built a tall tower deep in the woods. He built stairs that led up to a room at the very top, a room with one window. He and his wife took turns staying with the baby.

Rapunzel Story

But the witch had a magic ball. The ball showed her just where the baby was, at the top room of the tower. One day when the carpenter and his wife were both in the house, she cast a spell over both of them. They fell into a deep, deep sleep. And at once, the witch went to the tower.

At the top room, the witch said to the baby, I will call you Rapunzel. For that is the name of the lettuce that brought you to me. Now Rapunzel, you are mine!

But the witch did not know how to take care of a baby. Rapunzel grew into a child, and the witch did not even know how to cut her hair. The girl's blond hair grew longer and longer every day.

Rapunzel Story

All the witch could do was keep the child locked in the room at the very top of the tower. She told the girl that the world was a very bad place. That was why she could not leave the tower.

As she grew up, many times Rapunzel said to the witch, There is nothing here for me to do! Why must I stay in this tower all the time?

And the witch shouted back, I already told you so many times! The world is a very bad place. Now go comb your hair and be quiet.

But is it really so bad out there? Sometimes I hear people laughing down below, Rapunzel would say sometimes.

At such times the witch would yell, How many times do I have to repeat myself? Don't listen to anything you see or hear out there. The world is much worse than you think! You will stay in this tower forever, Rapunzel. So get used to it!

On her 12th birthday, Rapunzel said to the witch, I do not care what you say anymore! I am so tired of staying here alone all

the time! When you are gone, I will chip away at the door. I will make a hole. I will run down the stairs and outside, no matter what you say!

Think again! said the witch. With her power, she made all the stairs in the tower fall down. She made the doors close up. Now there was no way for Rapunzel to escape!

Chapter 3 A Singing Voice

By then, Rapunzel's hair had grown very, very long. Once the stairs were gone, when it was time for the witch to visit her in the tower, she would call from outside, Rapunzel, Rapunzel! Let down your hair!

Rapunzel would throw her long blond braid out of the window. The witch would grab hold of her hair like a rope. And that is how the witch climbed up the tower wall to the window in Rapunzel's room.

Rapunzel Story

Five more long years went by. Poor Rapunzel! She knew she must stay in the room. All she could do was to sing sad songs out of the window. Sometimes birds at the treetops would join in her songs. Then she would feel a bit better. But not much.

One day, a prince was riding through the woods. He heard a beautiful singing voice. Where was it coming from? He rode closer and closer to the sound. At last, he came to the tower.

This is odd! he said, looking around the tower wall. There is no door at the bottom. Yet someone is singing at the very top. How does anyone get in or out of there? Each day, the prince came back to the tower. There was something about that voice that pulled him back. Who was that young woman singing at the top? Could he ever meet her?

One day when the prince rode up, he saw an old woman standing below the tower. He jumped behind a tree to hide. It was a witch! He heard her call out, Rapunzel, Rapunzel! Let down your hair! A long blond braid was thrown out from a window at the very top. The old woman grabbed onto the braid. And she climbed the wall to the window at the top of the tower.

Ah, ha! said the prince. So that is how it is done! He waited.

Rapunzel Story

After a bit, the braid was thrown from the window again. The witch climbed back down the tower wall. Then she left.

The prince waited. He stepped up to the tower. In a voice that sounded as much like the witch as he could, he called out, Rapunzel, Rapunzel! Let down your hair! In a moment, the same long blond braid came out of the window. It worked! thought the prince. He climbed up the wall of the tower.

You can be sure that Rapunzel was very surprised to see the prince climb into her window. She had never seen a person up close before other than the witch, and never a man! Who are you? she said in fear.

Do not worry! said the prince. I am a friend.

But I do not know you, said Rapunzel.

I feel as if I know you, said the prince. Since I have heard you sing songs from up here day after day. You have a beautiful voice! And I love it when the birds sing with you, too.

Yes, I like that, too, said Rapunzel. It may be the only thing I like, since I must stay here in this same old tower, day after day, my whole life long. Rapunzel told the prince about the witch. She told him that since the world was such a very bad place, she must always stay in the tower room.

But the world is not as bad as she says! said the prince. He told Rapunzel about flowers and festivals, games and gardens. He told her about puppies and puddles, strawberries and secrets.

Many hours went by. At last, Rapunzel said he must go the witch may come back at any time! Very well, said the prince. But

I will be back tomorrow.

Rapunzel threw her braid out the window, and the prince climbed down. The next day, the prince climbed back up to Rapunzels room. He said, I have a surprise for you. He had brought strawberries for her.

As she tasted a strawberry Rapunzel thought, Now I know that what I was told is not true. The world can be a very fine place! I must get out of this tower as soon as I can. But how?

Rapunzel Story

Chapter 4 Plan to Escape

One day, the prince said, If only you could get out of this tower. I can come and go by climbing up the walls by holding onto your braid. But once I am down, how can you get down, too?

I know! said Rapunzel. Bring me a ball of silk each time you come. I can weave the silk into a ladder. Silk folds up so small the witch wont see it. When the ladder gets long enough to reach the ground, we will both be able to climb out of here.

Thats it! said the prince. Then he moved closer to Rapunzel. We will both be free. When we are out in the world, will you marry me?

Yes, said Rapunzel, I will. Every day after that, the prince brought a ball of silk to Rapunzel. Over time, she weaved the silk into a long ladder.

On Rapunzels 18th birthday the witch spoke to her in a sharp voice. Before you open your mouth this time, said the witch, know this. I am sick and tired of hearing you talk about how alone you are in the tower all the time. It isnt go to change, Rapunzel! Forever!

Who says Im alone in the room all the time? said Rapunzel.

What?! said the witch. Who has been up here with you?

No one! said Rapunzel at once, in fear. I mean, no one but you!

The witch did not believe her. She started to look everywhere in the room for something to prove that someone else had been there. Soon she found the ladder. She held it high in the air. She yelled, What is the meaning of this?

My friend the prince brought me the silk, said Rapunzel.

You will never see this prince again! yelled the witch. She took out a knife. Snip, snap, and Rapunzels lovely braid was cut off!

Holding the braid in one hand, the witch laughed an evil laugh. With a stroke of her magic, Rapunzel was cast away to a far-away desert. Then the witch stayed in the tower room. She knew that soon the prince would come back.

Rapunzel Story

Chapter 5 The Last Climb

The witch did not have to wait long. Soon the prince was calling at the bottom of the tower, in a voice that was supposed to sound like her own, Rapunzel, Rapunzel, let down your hair!

So that is how he did it! thought the witch. Holding tightly to one end of Rapunzels braid, she threw the braid out the window. The prince took hold and climbed up. When he got to the window, he was much surprised to see the witch!

Where is Rapunzel? he called out. What have you done with her?

You will never see your Rapunzel again! yelled the witch.

The witch pushed the prince so hard that he lost hold of the window. Down, down, he fell! The prince landed on some bushes below. That helped with the fall, but the bushes had sharp thorns. Some of the thorns went into his eyes. The prince was blind!

Chapter 6 The Desert

For two years the poor blind prince wandered the world, looking for Rapunzel. From morning to night he called for her, but it was no use. At last, he reached a desert. One day, he heard a beautiful voice singing. Oh! he thought. I know that voice! It was his dear Rapunzel! He went closer and closer to the voice he knew so well.

My prince! called Rapunzel when she saw him. The two of them hugged tight. Two tears of joy fell into the eyes of the prince. All at once, he could see again!

Rapunzel Story

And what happened next, well, I'm sure you can guess! The prince and Rapunzel went back to the kingdom where the prince lived. They were married as soon as they could. The prince became king of the land and Rapunzel became queen. The two of them lived happily ever after.

C.1.4 The Frog Prince: The Story of the Princess and the Frog

Source: <https://www.storiestogrowby.org/story/princess-and-the-frog-story-bedtime-stories-for-kids/>

Once upon a time there was a Princess. Many a suitor came to the palace to win her hand in marriage, but it seemed to the Princess that each one of them looked at her without really seeing her at all.

They act like there's nothing more to a princess than her fine crown and royal dresses, she said to herself with a frown.

One afternoon after one of these visits, the Princess thought, Sometimes I wish I were little again. She found her favorite ball from childhood, the one that sparkled when she threw it up high to the sun. She took the ball to the palace yard and threw it higher and higher. One time she threw it extra high and when she ran to catch the ball, she tripped on a tree stump. The ball fell and plopped right down into the royal well! She raced over to fetch her ball before it dropped too far, but by the time she got there she could no longer see it in the water.

The ball fell and plopped down into the royal well!

Oh no! she moaned, This is terrible! Just then a small green frog poked its head above the water.

The Princess and the Frog

Maybe I can help you, said the Frog.

Yes, said the Princess. Please get my ball!

No problem, said the Frog. But first there's something I must ask of you.

What do you mean? said the Princess.

It's for you to spend time with me today, said the Frog.

I'm not sure I know what that means, said the Princess.

Just spend time with me today, repeated the Frog.

All right then, fine! said the Princess. Now please, get my ball!

I'm on it, said the Frog. He dived deep into the well. A few moments later, up he came with the ball held high in one hand.

Thank you, said the Princess, taking it from him.

The Princess and the Frog Story

She turned to go.

Wait a minute! said the Frog. You promised to spend time with me today!

I already did, she said with a shrug. And the Princess walked back to the palace.

That night at dinner with her family and the royal advisers, there was a knock on the door. The servant opened the door and saw no one there. The Frog, standing down low, cleared his throat. The Princess promised to spend time with me today, said the Frog in as loud a voice as he could. So here I am.

That night at dinner with her family and the royal advisers, there was a knock on the door.

Daughter! said the King from the far end of the table. Did you promise to spend time with this Frog, as he claims?

Sort of, said the Princess. After a pause, she added, Oh very well, come on in.

The servants quickly set a new place setting for the Frog, and he hopped over to the royal dining table.

Princess and the Frog

Conversation turned to a topic of concern in the kingdom. None of the royal advisers knew what to do.

Father, if I may, said the Princess. Perhaps we could

Stop! said the King, cutting her off. I have enough advisers, believe me.

If I may, said the Frog, and it was the first time he had spoken at the table. There's more to a princess than her fine crown and royal dresses.

The Princess stared at the Frog. How could this little frog more than anyone else understand such a thing?

If I may, said the Frog, and it was the first time he has spoken at the table.

After dinner, the Frog bowed to the Princess. He said, You have done what you said you would do. I suppose it's time now for me to go.

No wait! said the Princess, it's not that late. How about a walk in the garden?

The Frog was delighted. The two of them walked in the royal garden, the Frog hopping along the stone wall so he and the Princess were at the same level and could talk easily. They laughed about many things. Later, when the sun set, they admired the deep rosy reds it cast in the sky.

The Princess said, You know, being with you tonight was a lot more fun than I thought.

I had a very good time, too, said the Frog.

Who knew? said the Princess with a laugh. She leaned over and kissed the Frog lightly on his cheek.

The Princess and the Frog

At once, there was a puff of clouds and smoke. The small green frog had changed into a young prince! The Princess jumped back in surprise, and who could blame her? The Prince quickly told her not to worry, that all was well. Years before, an evil witch had put a spell on him that he must stay a frog until he was kissed by a princess. The witch had laughed an evil laugh, saying, Like THAT will ever happen! But it did!

Now the Prince and Princess could get to know each other better. Years later, after they were married, they had a beautiful setting made for the ball and placed it on their royal dining table. And when the sunlight shone in through the palace windows, the ball sparkled for all to see.

C.1.5 Aladdin and the Magic Lamp from The Arabian Nights

Category: Children Story

Source: <https://www.storiestogrowby.org/story/aladdin-story-from-the-arabian-nights-bedtime-stories-folk-tales-for-kids/>

Once upon a time, a young mans father died. Aladdin, as that was the young mans name, took his fathers place in running the family store with his mother. One day, a stranger walked into the store.

I am your uncle, said the stranger to Aladdin. I have come to see you.

But my father never spoke of a brother, said Aladdin.

Aladdins mother turned around. My husband had no brother, said she to the stranger, narrowing her eyes.

I assure you it is true, said the stranger. Years ago your husband and I agreed that if something should happen to him, since I have been very fortunate in my life, I would help to bring the same good fortune to your family.

The mother was interested. What do you have in mind? she said.

I know of a secret place that holds many riches, said the stranger. I will take your son. With the wealth he will find there, you and he will be set for life.

Aladdin Story Cave of Wonders

And so the mother agreed. The old man and the boy traveled for days throughout the desert. At last they came to a cave. You must know that I learned a bit of magic in my life, said the old man to Aladdin. Dont be surprised by anything you might see.

They stepped inside the cave. Pitch-black it was. The old man shook open his fist and a ball of light suddenly appeared, brightening the cave. Under the light with one long finger, he drew the shape of a circle over the ground. He pulled from his pocket some red dust and threw it over the circle, and at the same time said some magic words. The earth trembled a little before them. The floor of the cave cracked open, and the cracks grew wider and deeper. Then up from below the ground rose a giant white quartz crystal and it filled the circle.

Do not be alarmed, said the magician. Under this giant white crystal lies a treasure that is to be yours.

He chanted a few magic words and the giant crystal rose up several feet in the air, moved to the side and landed. Aladdin peered into the hole. He saw steps that led down to a dark hole.

Fear nothing, said the magician to Aladdin. But obey me. Go down, and at the foot of the steps, follow a long hall. You will walk through a garden of fruit trees. You must touch nothing of them. Walk on till you come to a large flat stone and on the stone will be a lighted lamp. Pour out the oil in the lamp and bring it to me. Now go!

Aladdin slowly stepped down the stairs. Through the garden of fruit trees and marvelous to behold, the trees held fruits that sparkled and shone. He could not help but reach out and touch one.

Then too late he remembered what his uncle had said. But nothing terrible happened. So he figured he might as well put the fine jewel-fruit in his vest pocket. Then he plucked another and another jewel-fruit, till all his pockets were filled.

Aladdin came to the large flat stone, and on it was a lighted lamp, just as his uncle had said. He poured out the oil and took it back to the opening of the cave.

Aladdin came to the large flat stone, and on it was a lighted lamp, just as his uncle had said.

Aladdin called out, Here I am, Uncle!

The magician called out in a great hurry, Give me the lamp!

Just as soon as I'm up, said Aladdin, wondering why the magician seemed in such a hurry.

No, give me the lamp NOW! cried the old man, reaching down his hand. For you see, the only way the lamp could come out of the cave was as a gift, from one person to another.

The magician knew this, and he wanted to get the lamp from the boy as soon as he could, and then kill him. Aladdin felt a chill in the air. Something was wrong. Somehow he knew he must not give up that lamp.

Let me up first," said Aladdin. "Then will I give you the lamp.

Aladdin felt a chill in the air. Something was wrong.

The magician was furious. He fell into a rage and barked out more magical words. The giant white quartz crystal rose up, hovered over the hole and landed. All went dark below. Aladdin was trapped!

For two days, Aladdin despaired. Why didn't I just hand over this old lamp? Who cares about it, anyway? Whatever might have come of it, it couldn't have been worse than this! What was I thinking?

Rubbing the lamp, he moaned, Oh, how I wish I could get out of here!

At once, a huge Genie rose up into the air. You are my master! boomed the Genie. Was that your first wish - to get out of this cave? Three wishes are yours to command.

Genie of the Magic Lamp

Aladdin's mouth fell open, amazed. He mumbled yes, of course! More than anything he wanted to get out of the cave and go home! The very next moment, Aladdin was outside his own home, still holding the lamp and with all his jewel-fruits in his vest pockets.

His mother could not believe the tale her son told her. Magic lamp? she laughed. That old thing? She took the lamp, grabbed a rag, and started to clean it. If there were really a Genie in this old lamp, I would say to it, Genie, make a feast for my son and me, and serve it on plates of gold!

You can imagine the mother's surprise! The Genie rose up out of the lamp, and a feast fit for a king weighed down her kitchen table, on plates of glimmering gold.

You can imagine the mother's surprise!

Mother and son enjoyed a feast like no other. Then the mother washed and sold the gold plates, and bought necessary things to live. From then on, Aladdin and his mother lived well.

One day, Aladdin thought to himself, Why think small? With my jewel-fruits, I could marry the princess and become the prince of this land!

His mother laughed. You can't just go to a palace with some fine gifts and expect to marry the princess! But Aladdin urged her to try. They wrapped some of the jewel-fruits in silk cloth, and the mother went to the palace.

The guards stopped her at once. But as she insisted she had something very valuable for the Sultan, they let her in.

They wrapped some of the jewel-fruits in silk cloth.

Said the Sultan, What have you brought me in those silk rags?

She showed him the jewel-fruits. The Sultan was impressed. But if your son is as worthy of my daughter as you say, he must bring me 40 golden trays of the same gems, carried in by servants.

The mother went home and told her son the Sultans demand. It's no problem, said Aladdin. Call for the Genie and make your second wish. And so his mother rubbed the lamp and made her second wish. Before long, she was at the steps of the Sultans palace with 40 golden trays of the jewel-fruits, carried in by as many servants.

The Sultan was pleased. But you cannot think this is enough to win the hand of my daughter! he said. To truly win my favor, your son must build a golden palace for he and my daughter to live.

The Mother brought back this news, too. So for her third wish, the Mother asked the Genie to create a golden palace. The next morning, right outside the Sultans bedroom, appeared a huge golden palace, gleaming in the sun.

Aladdin and the Magic Lamp

Meanwhile back at Aladdins home, his Mother said, It is time for you to go, my son, to meet your princess. Her wishes spent, she gave him the lamp.

The next morning, the Sultan called for his daughter. Look at this palace! he said, pointing out the window. This is the husband for you!

What do you mean, Father? said his daughter. What do you know about this man? Have you ever met him?

Whats there to know? said the Sultan. He can make a golden palace appear overnight. Hes even more powerful than my royal adviser, the Vizier.

Whats there to know? said the Sultan. "He can make a golden palace appear overnight."

Yesterday, your Vizier was most powerful man in the kingdom," said his daughter, "and I was to marry him. Today, this stranger is the most powerful one and Im to marry him. Why do you think it matters to me whos the most powerful?

Aladdin Princess Jasmine

It matters to ME! said the Sultan. In a lower voice he said, Daughter, youre just excited to get such a fine husband.

I cant believe this! The princess threw her arms up in despair, and she left.

In her dressing room, the princess groaned. To Nadia, her lady-in-waiting, she said, My father is determined to marry me off, no matter what!

But Madam, said Nadia, isnt this wonderful stranger an excellent match for you?

The princess sighed. She looked at her lady-in-waiting. You dont know how lucky you are, she said. I would rather live your life than be handed off in this way.

And I would rather live yours, said Nadia. The two of them stared at each other for a couple of moments. They were about the same height, with the same color hair. With all the scarves maidens like them wore

Lets do it! they said together. And the two of them changed clothes.

Just then, Aladdin was riding to the Sultans palace on a white horse, ready to meet his bride. The Sultan warmly greeted him.

Lets do it! they said together.

Stay here in my palace until the preparations for your wedding are complete, he said. Aladdin could not meet the princess until their wedding day. He caught a glimpse of Nadia from a distance, covered in scarves, thinking she was the true princess. Aladdin, the Sultan, and everyone else in the palace waited with growing excitement for the wedding day.

Except for one person. The uncle-magician who had left Aladdin trapped in the cave was also the Sultans Vizier.

Aladdin and the Magic Lamp

He had recognized Aladdin at once. He knew there could be only one reason the young man could present all this magic to the Sultan. Aladdin must have escaped from the cave, and with the lamp!

I will get my revenge! swore the Vizier. If anyone is to have the lamp, it is ME! By his magic, he could tell where Aladdin had hidden the lamp. While Aladdin was sleeping, the Vizier crept in and took it.

If anyone is to have the lamp, it is ME!"

In a quiet place, the Vizier made his first wish: Genie, do as I say. I want you to take Aladdins palace to a faraway place in the desert that no one can find!

What the Vizier did not know was at that very moment, Nadia was exploring Aladdins palace. And there is something else the Vizier did not know. The Genie thought the Vizier had commanded to be taken away also, along with the palace. So the Genie sent the Vizier, the golden palace, and Nadia inside it, all together to the faraway place in the desert.

The next morning, the Sultan awoke and saw nothing outside his bedroom window where Aladdins palace had stood the day before. The next moment his servants rushed in, announcing that the princess had disappeared. Furious, he called for Aladdin.

What have you done? he yelled in a rage. Because of your magic tricks I have lost my daughter! You must bring her back to me in three days or it will cost you your head!

"What have you done?" he yelled in a rage.

Aladdin thought he would simply use his second wish and the Genie would bring back the princess and the castle too. But his magic lamp was gone - he looked everywhere!

In despair, Aladdin could do nothing but to leave the Sultans palace on the white horse he had rode in on. Sadly, he rode from town to town but no one knew anything about a palace that had appeared overnight, no to mention one with a princess inside.

You may wonder, where was the true princess all this time? Dressed as a servant girl, she had crept out of the palace the very day she had switched clothes with Nadia. Down to the marketplace she had gone, and there she met an aging merchant. The old merchant told her he was tired from riding so many years from town to town, selling his potions and perfumes.

The princess was dressed humbly, yet she still carried herself like royalty. She gained the confidence of the old merchant, and when she offered to ride his camel train for him and share what she earned, he was delighted. That is how our princess found herself up clop-clopping through the desert, selling potions and perfumes from town to town.

Aladdin and the Magic Lamp

Two days passed. Aladdin was no closer to finding his lost palace than he had been before he left the Sultan. Crouched in front of his tent, Aladdin held his head in his hands.

Why the sad face? The princess was riding by and she stopped her camel train. Perhaps a potion will make you feel better.

"Why the sad face?"

No, thank you, said Aladdin. The only thing that could help is if I could bring back a princess and find my lost palace. You see, my palace vanished overnight to a place I know not where. The princess was probably inside it. Oh, this is an impossible task!

Maybe not, said the princess. In my travels, I heard of a palace in the desert that appeared out of nowhere, not long ago.

Really? said Aladdin. He looked up. Do you know where?

I think so. I could take you there. If we left now, we could get there by morning.

Id be so grateful! said Aladdin. He had left all the jewel-fruits with his Mother except one. This he offered to the camel-rider as payment.

Oh, keep it, said she with a wave of her hand. Its no trouble. Bring your horse to ride alongside my camel.

Really? said Aladdin. He looked up. Do you know where?

Riding through the night, the two of them spoke of many things. Aladdin marveled at the young ladys easy manner and generous spirit. He somehow knew she could be trusted. Before long, he told her his story of how he had discovered the magic lamp in the cave and how it had been stolen from him, along with the palace.

As the mornings light brightened, they were riding between two very tall walls of rock, rose-colored they were, with thin bands of white and blue. Suddenly the rock walls ended, and they arrived at a clearing.

Look! said the princess, pointing ahead. Is that it?

Aladdin and the Magic Lamp

It is! Aladdin cried out with joy, recognizing his palace. I hope the princess is still in there!" he said. "Though without my lamp, I have no way to get them both back in time.

Just then Nadia, who had been carried away along with the palace as you no doubt remember, was looking out the window at the new guests. To her surprise, she recognized the rider of the camel train as none other than her beloved former mistress. She waved at them both to come to the front door.

The servants let in the guests. Nadia took them to the drawing room and shut the door. She said, Mistress! How glad I am to see you!

I'm glad to see you too, Nadia.

Aladdin was amazed. You two know each other?

But the princess only said to Nadia, Tell me, how do you find being a princess?

Aladdin was amazed. You two know each other?

At first, the gowns were marvelous, she said. Everything I dreamed of! And I liked well enough all the attention I got. But when I was carried away with this palace, the Vizier came with it, too. For the last two days he has done nothing but fly about in a rage and smash things. He locked me up in here!

Thats terrible! said the princess.

Theres more, said Nadia. He said with his lamp, that tomorrow well return to the Sultans land and I will have to marry him!

He saidwith his lamp? Aladdin and the princess looked at each other.

The princess turned to Nadia. Wait a minute! I have a plan.

The princess turned to Nadia. Wait a minute! I have a plan.

The princess gave Nadia one of the sleeping potions in her stock. She told Nadia that when the Vizier returned that night, she must pour the sleeping potion into his wine. He would fall into a sleep so deep that he would not be awakened by any noise. That is what she did. When the wicked man was snoring, Nadia, the princess, and Aladdin searched everywhere for the magic lamp. At last they found it!

The lamp in his hands again, Aladdin said, Now I can make a second wish. I am going to wish for this castle and everyone in it to go back to the Sultans kingdom except for the Vizier.

Wait! said the princess. "Leave me behind, too.

Aladdin urged her to come with him, but the princess would have none of it. She liked too well the life of freedom she led.

Aladdin did not like at all that she would be left behind with the Vizier. But she assured him the Vizier would not awaken for hours, and she would have plenty of time to get far away.

So Aladdin rubbed the lamp and stated his wish to the Genie.

Aladdin and the Magic Lamp

In a whoosh, Aladdin, the palace and Nadia were all transported back to the very spot where the palace had stood before.

The Sultan was delighted to have his daughter back, or you might say, the young woman he believed to be his daughter, covered as she was in scarves. We will hold the wedding in three days! the Sultan said to Aladdin.

Yet a sadness was growing in Aladdins heart. Nadia was indeed a nice young woman, and pleasant to look at, too. But there was something about that woman who rode the camel train, selling perfumes and potions. He could not get out of his mind the sound of her laugh, her clever mind, and the comfort of her company. At last, he rubbed the lamp.

Master, said the Genie, Is it mountains of jewels you want for your third wish, power over all the neighboring lands, or the strength of 100 men?

But there was something about that woman who rode the camel train.

None of that, said Aladdin. I wish you to take me to that young woman I met, the camel rider, the seller of perfumes and potions.

But Master, this is your third and last wish! said the Genie. What if you were to offer this woman your heart, and she didnt choose you back? Youll miss your chance to marry the Sultans daughter and become a prince.

I dont care! said Aladdin. I must share with her what is in my heart. Whatever comes of it, so be it.

But Master, this is your third and last wish! said the Genie.

So Aladdin made his third and last wish and was taken to the true princess. In her travels, she was not all that far from the Sultans land, as it turns out. Aladdin shared his true feelings to her and she returned the same feelings.

Aladdin and Princess Jasmine

She told him her story that she had been born a princess but now was happier living as a traveling merchant. Aladdin said he wanted nothing better than to spend the rest of his days with her by his side. And so they agreed to marry and together ride the camel train, selling potions and perfumes from town to town.

Then such surprising news! Aladdin and the princess learned that the Sultan had suddenly died. Said Aladdin to his new bride, Since your father is gone, would you return now to your fathers palace? We could rule the kingdom together, side by side.

And so Aladdin and the princess returned to the palace. Nadia was very pleased to see them. She gladly stepped down to serve again as lady-in-waiting to the princess. For the rest of their lives, Aladdin and the princess ruled well and lived in happiness, as you should, too.

C.2 Modern Ghost Stories

C.2.1 The Shadows on The Wall

Source: <http://www.gutenberg.org/files/15143/15143-h/15143-h.htm#Shadows>

"Henry had words with Edward in the study the night before Edward died," said Caroline Glynn.

She spoke not with acrimony, but with grave severity. Rebecca Ann Glynn gasped by way of assent. She sat in a wide flounce of black silk in the corner of the sofa, and rolled terrified eyes from her sister Caroline to her sister Mrs. Stephen Brigham, who had been Emma Glynn, the one beauty of the family. The latter was beautiful still, with a large, splendid, full-blown beauty, she filled a great rocking-chair with her superb bulk of femininity, and swayed gently back and forth, her black silks whispering and her black frills fluttering. Even the shock of death for her brother Edward lay dead in the house could not disturb her outward serenity of demeanor.

But even her expression of masterly placidity changed before her sister Caroline's announcement and her sister Rebecca Ann's gasp of terror and distress in response.

"I think Henry might have controlled his temper, when poor Edward was so near his end," she said with an asperity which disturbed slightly the roseate curves of her beautiful mouth.

"Of course he did not know," murmured Rebecca Ann in a faint tone.

"Of course he did not know it," said Caroline quickly. She turned on her sister with a strange, sharp look of suspicion. Then she shrank as if from the other's possible answer.

Rebecca gasped again. The married sister, Mrs. Emma Brigham, was now sitting up straight in her chair; she had ceased rocking, and was eyeing them both intently with a sudden accentuation of family likeness in her face.

"What do you mean?" said she impartially to them both. Then she, too, seemed to shrink before a possible answer. She even laughed an evasive sort of laugh.

"Nobody means anything," said Caroline firmly. She rose and crossed the room toward the door with grim decisiveness.

"Where are you going?" asked Mrs. Brigham.

"I have something to see to," replied Caroline, and the others at once knew by her tone that she had some solemn and sad duty to perform in the chamber of death.

"Oh," said Mrs. Brigham.

After the door had closed behind Caroline, she turned to Rebecca.

"Did Henry have many words with him?" she asked.

"They were talking very loud," replied Rebecca evasively.

Mrs. Brigham looked at her. She had not resumed rocking. She still sat up straight, with a slight knitting of intensity on her fair forehead, between the pretty rippling curves of her auburn hair.

"Did you ever hear anything?" she asked in a low voice with a glance toward the door.

"I was just across the hall in the south parlor, and that door was open and this door ajar," replied Rebecca with a slight flush.

"Then you must have"

"I couldn't help it."

"Everything?"

"Most of it."

"What was it?"

"The old story."

"I suppose Henry was mad, as he always was, because Edward was living on here for nothing, when he had wasted all the money father left him."

Rebecca nodded, with a fearful glance at the door.

When Emma spoke again her voice was still more hushed. "I know how he felt," said she. "It must have looked to him as if Edward was living at his expense, but he wasn't."

"No, he wasn't."

"And Edward had a right here according to the terms of father's will, and Henry ought to have remembered it."

"Yes, he ought."

"Did he say hard things?"

"Pretty hard, from what I heard."

"What?"

"I heard him tell Edward that he had no business here at all, and he thought he had better go away."

"What did Edward say?"

"That he would stay here as long as he lived and afterward, too, if he was a mind to, and he would like to see Henry get him out; and then"

"What?"

"Then he laughed."

"What did Henry say?"

"I didn't hear him say anything, but"

"But what?"

"I saw him when he came out of this room."

"He looked mad?"

"You've seen him when he looked so."

Emma nodded. The expression of horror on her face had deepened.

"Do you remember that time he killed the cat because she had scratched him?"

"Yes. Don't!"

Then Caroline reentered the room; she went up to the stove, in which a wood fire was burning it was a cold, gloomy day of fall and she warmed her hands, which were reddened from recent washing in cold water.

Mrs. Brigham looked at her and hesitated. She glanced at the door, which was still ajar; it did not easily shut, being still swollen with the damp weather of the summer. She rose and pushed it together with a sharp thud, which jarred the house. Rebecca started painfully with a half-exclamation. Caroline looked at her disapprovingly.

"It is time you controlled your nerves, Rebecca," she said.

Mrs. Brigham, returning from the closed door, said imperiously that it ought to be fixed, it shut so hard.

"It will shrink enough after we have had the fire a few days," replied Caroline.

"I think Henry ought to be ashamed of himself for talking as he did to Edward," said Mrs. Brigham abruptly, but in an almost inaudible voice.

"Hush," said Caroline, with a glance of actual fear at the closed door.

"Nobody can hear with the door shut. I say again I think Henry ought to be ashamed of himself. I shouldn't think he'd ever get over it, having words with poor Edward the very night before he died. Edward was enough sight better disposition than Henry, with all his faults."

"I never heard him speak a cross word, unless he spoke cross to Henry that last night. I don't know but he did from what Rebecca overheard."

"Not so much cross, as sort of soft, and sweet, and aggravating," sniffed Rebecca.

"What do you really think ailed Edward?" asked Emma in hardly more than a whisper. She did not look at her sister.

"I know you said that he had terrible pains in his stomach, and had spasms, but what do you think made him have them?"

"Henry called it gastric trouble. You know Edward has always had dyspepsia."

Mrs. Brigham hesitated a moment. "Was there any talk of an examination?" said she.

Then Caroline turned on her fiercely.

"No," said she in a terrible voice. "No."

The three sisters' souls seemed to meet on one common ground of terrified understanding through their eyes.

The old-fashioned latch of the door was heard to rattle, and a push from without made the door shake ineffectually. "It's Henry," Rebecca sighed rather than whispered. Mrs. Brigham settled herself, after a noiseless rush across the floor, into her rocking-chair again, and was swaying back and forth with her head comfortably leaning back, when the door at last yielded and Henry Glynn entered. He cast a covertly sharp, comprehensive glance at Mrs. Brigham with her elaborate calm; at Rebecca quietly huddled in the corner of the sofa with her handkerchief to her face and only one small uncovered reddened ear as attentive as a dog's, and at Caroline sitting with a strained composure in her armchair by the stove. She met his eyes quite firmly with a look of inscrutable fear, and defiance of the fear and of him.

Henry Glynn looked more like this sister than the others. Both had the same hard delicacy of form and aquilinity of feature. They confronted each other with the pitiless immovability of two statues in whose marble lineaments emotions were fixed for all eternity.

Then Henry Glynn smiled and the smile transformed his face. He looked suddenly years younger, and an almost boyish recklessness appeared in his face. He flung himself into a chair with a gesture which was bewildering from its incongruity with his general appearance. He leaned his head back, flung one leg over the other, and looked laughingly at Mrs. Brigham.

"I declare, Emma, you grow younger every year," he said.

She flushed a little, and her placid mouth widened at the corners. She was susceptible to praise.

"Our thoughts to-day ought to belong to the one of us who will never grow older," said Caroline in a hard voice.

Henry looked at her, still smiling. "Of course, we none of us forget that," said he, in a deep, gentle voice; "but we have to speak to the living, Caroline, and I have not seen Emma for a long time, and the living are as dear as the dead."

"Not to me," said Caroline.

She rose and went abruptly out of the room again. Rebecca also rose and hurried after her, sobbing loudly.

Henry looked slowly after them.

"Caroline is completely unstrung," said he.

Mrs. Brigham rocked. A confidence in him inspired by his manner was stealing over her. Out of that confidence she spoke quite easily and naturally.

"His death was very sudden," said she.

Henry's eyelids quivered slightly but his gaze was unswerving.

"Yes," said he, "it was very sudden. He was sick only a few hours."

"What did you call it?"

"Gastric."

"You did not think of an examination?"

"There was no need. I am perfectly certain as to the cause of his death."

Suddenly Mrs. Brigham felt a creep as of some live horror over her very soul. Her flesh prickled with cold, before an inflection of his voice. She rose, tottering on weak knees.

"Where are you going?" asked Henry in a strange, breathless voice.

Mrs. Brigham said something incoherent about some sewing which she had to do some black for the funeral and was out of the room. She went up to the front chamber which she occupied. Caroline was there. She went close to her and took her hands, and the two sisters looked at each other.

"Don't speak, don't, I won't have it!" said Caroline finally in an awful whisper.

"I won't," replied Emma.

That afternoon the three sisters were in the study.

Mrs. Brigham was hemming some black material. At last she laid her work on her lap.

"It's no use, I cannot see to sew another stitch until we have a light," said she.

Caroline, who was writing some letters at the table, turned to Rebecca, in her usual place on the sofa.

"Rebecca, you had better get a lamp," she said.

Rebecca started up; even in the dusk her face showed her agitation.

"It doesn't seem to me that we need a lamp quite yet," she said in a piteous, pleading voice like a child's.

"Yes, we do," returned Mrs. Brigham peremptorily. "I can't see to sew another stitch."

Rebecca rose and left the room. Presently she entered with a lamp. She set it on the table, an old-fashioned card-table which was placed against the opposite wall from the window. That opposite wall was taken up with three doors; the one small space was occupied by the table.

"What have you put that lamp over there for?" asked Mrs. Brigham, with more of impatience than her voice usually revealed. "Why didn't you set it in the hall, and have done with it? Neither Caroline nor I can see if it is on that table."

"I thought perhaps you would move," replied Rebecca hoarsely.

"If I do move, we can't both sit at that table. Caroline has her paper all spread around. Why don't you set the lamp on the study table in the middle of the room, then we can both see?"

Rebecca hesitated. Her face was very pale. She looked with an appeal that was fairly agonizing at her sister Caroline.

"Why don't you put the lamp on this table, as she says?" asked Caroline, almost fiercely. "Why do you act so, Rebecca?"

Rebecca took the lamp and set it on the table in the middle of the room without another word. Then she seated herself on the sofa and placed a hand over her eyes as if to shade them, and remained so.

"Does the light hurt your eyes, and is that the reason why you didn't want the lamp?" asked Mrs. Brigham kindly.

"I always like to sit in the dark," replied Rebecca chokingly. Then she snatched her handkerchief hastily from her pocket and began to weep. Caroline continued to write, Mrs. Brigham to sew.

Suddenly Mrs. Brigham as she sewed glanced at the opposite wall. The glance became a steady stare. She looked intently, her work suspended in her hands. Then she looked away again and took a few more stitches, then she looked again, and again turned to her task. At last she laid her work in her lap and stared concentratedly. She looked from the wall round the room, taking note of the various objects. Then she turned to her sisters.

"What is that?" said she.

"What?" asked Caroline harshly.

"That strange shadow on the wall," replied Mrs. Brigham.

Rebecca sat with her face hidden; Caroline dipped her pen in the inkstand.

"Why don't you turn around and look?" asked Mrs. Brigham in a wondering and somewhat aggrieved way.

"I am in a hurry to finish this letter," replied Caroline shortly.

Mrs. Brigham rose, her work slipping to the floor, and began walking round the room, moving various articles of furniture, with her eyes on the shadow.

Then suddenly she shrieked out:

"Look at this awful shadow! What is it? Caroline, look, look! Rebecca, look! What is it?"

All Mrs. Brigham's triumphant placidity was gone. Her handsome face was livid with horror. She stood stiffly pointing at the shadow.

Then after a shuddering glance at the wall Rebecca burst out in a wild wail.

"Oh, Caroline, there it is again, there it is again!"

"Caroline Glynn, you look!" said Mrs. Brigham. "Look! What is that dreadful shadow?"

Caroline rose, turned, and stood confronting the wall.

"How should I know?" she said.

"It has been there every night since he died!" cried Rebecca.

"Every night?"

"Yes; he died Thursday and this is Saturday; that makes three nights," said Caroline rigidly. She stood as if holding her calm with a vise of concentrated will.

"It it looks likelike" stammered Mrs. Brigham in a tone of intense horror.

"I know what it looks like well enough," said Caroline. "I've got eyes in my head."

"It looks like Edward," burst out Rebecca in a sort of frenzy of fear. "Only"

"Yes, it does," assented Mrs. Brigham, whose horror-stricken tone matched her sisters', "only Oh, it is awful! What is it, Caroline?"

"I ask you again, how should I know?" replied Caroline. "I see it there like you. How should I know any more than you?"

"It must be something in the room," said Mrs. Brigham, staring wildly around.

"We moved everything in the room the first night it came," said Rebecca; "it is not anything in the room."

Caroline turned upon her with a sort of fury. "Of course it is something in the room," said she. "How you act! What do you mean talking so? Of course it is something in the room."

"Of course it is," agreed Mrs. Brigham, looking at Caroline suspiciously. "It must be something in the room."

"It is not anything in the room," repeated Rebecca with obstinate horror.

The door opened suddenly and Henry Glynn entered. He began to speak, then his eyes followed the direction of the others. He stood staring at the shadow on the wall.

"What is that?" he demanded in a strange voice.

"It must be due to something in the room," Mrs. Brigham said faintly.

Henry Glynn stood and stared a moment longer. His face showed a gamut of emotions. Horror, conviction, then furious incredulity. Suddenly he began hastening hither and thither about the room. He moved the furniture with fierce jerks, turning ever to see the effect upon the shadow on the wall. Not a line of its terrible outlines wavered.

"It must be something in the room!" he declared in a voice which seemed to snap like a lash.

His face changed, the inmost secrecy of his nature seemed evident upon his face, until one almost lost sight of his lineaments. Rebecca stood close to her sofa, regarding him with woeful, fascinated eyes. Mrs. Brigham clutched Caroline's hand. They both stood in a corner out of his way. For a few moments he raged about the room like a caged wild animal. He moved every piece of furniture; when the moving of a piece did not affect the shadow he flung it to the floor.

Then suddenly he desisted. He laughed.

"What an absurdity," he said easily. "Such a to-do about a shadow."

"That's so," assented Mrs. Brigham, in a scared voice which she tried to make natural. As she spoke she lifted a chair near her.

"I think you have broken the chair that Edward was fond of," said Caroline.

Terror and wrath were struggling for expression on her face. Her mouth was set, her eyes shrinking. Henry lifted the chair with a show of anxiety.

"Just as good as ever," he said pleasantly. He laughed again, looking at his sisters. "Did I scare you?" he said. "I should think you might be used to me by this time. You know my way of wanting to leap to the bottom of a mystery, and that shadow does look queer, like and I thought if there was any way of accounting for it I would like to without any delay."

"You don't seem to have succeeded," remarked Caroline dryly, with a slight glance at the wall.

Henry's eyes followed hers and he quivered perceptibly.

"Oh, there is no accounting for shadows," he said, and he laughed again. "A man is a fool to try to account for shadows."

Then the supper bell rang, and they all left the room, but Henry kept his back to the wall as did, indeed, the others.

Henry led the way with an alert motion like a boy; Rebecca brought up the rear. She could scarcely walk, her knees trembled so.

"I can't sit in that room again this evening," she whispered to Caroline after supper.

"Very well; we will sit in the south room," replied Caroline. "I think we will sit in the south parlor," she said aloud; "it isn't as damp as the study, and I have a cold."

So they all sat in the south room with their sewing. Henry read the newspaper, his chair drawn close to the lamp on the table. About nine o'clock he rose abruptly and crossed the hall to the study. The three sisters looked at one another. Mrs. Brigham rose, folded her rustling skirts compactly round her, and began tiptoeing toward the door.

"What are you going to do?" inquired Rebecca agitatedly.

"I am going to see what he is about," replied Mrs. Brigham cautiously.

As she spoke she pointed to the study door across the hall; it was ajar. Henry had striven to pull it together behind him, but it had somehow swollen beyond the limit with curious speed. It was still ajar and a streak of light showed from top to bottom.

Mrs. Brigham folded her skirts so tightly that her bulk with its swelling curves was revealed in a black silk sheath, and she went with a slow toddle across the hall to the study door. She stood there, her eye at the crack.

In the south room Rebecca stopped sewing and sat watching with dilated eyes. Caroline sewed steadily. What Mrs. Brigham, standing at the crack in the study door, saw was this:

Henry Glynn, evidently reasoning that the source of the strange shadow must be between the table on which the lamp stood and the wall, was making systematic passes and thrusts with an old sword which had belonged to his father all over and through the intervening space. Not an inch was left unpierced. He seemed to have divided the space into mathematical sections. He brandished the sword with a sort of cold fury and calculation; the blade gave out flashes of light, the shadow remained unmoved. Mrs. Brigham, watching, felt herself cold with horror.

Finally Henry ceased and stood with the sword in hand and raised as if to strike, surveying the shadow on the wall threateningly. Mrs. Brigham toddled back across the hall and shut the south room door behind her before she related what she had seen.

"He looked like a demon," she said again. "Have you got any of that old wine in the house, Caroline? I don't feel as if I could stand much more."

"Yes, there's plenty," said Caroline; "you can have some when you go to bed."

"I think we had all better take some," said Mrs. Brigham. "Oh, Caroline, what"

"Don't ask; don't speak," said Caroline.

"No, I'm not going to," replied Mrs. Brigham; "but"

Soon the three sisters went to their chambers and the south parlor was deserted. Caroline called to Henry in the study to put out the light before he came upstairs. They had been gone about an hour when he came into the room bringing the lamp which had stood in the study. He set it on the table, and waited a few minutes, pacing up and down. His face was terrible, his fair complexion showed livid, and his blue eyes seemed dark blanks of awful reflections.

Then he took up the lamp and returned to the library. He set the lamp on the center table and the shadow sprang out on the wall. Again he studied the furniture and moved it about, but deliberately, with none of his former frenzy. Nothing affected the shadow. Then he returned to the south room with the lamp and again waited. Again he returned to the study and placed the lamp on the table, and the shadow sprang out upon the wall. It was midnight before he went upstairs. Mrs. Brigham and the other sisters, who could not sleep, heard him.

The next day was the funeral. That evening the family sat in the south room. Some relatives were with them. Nobody entered the study until Henry carried a lamp in there after the others had retired for the night. He saw again the shadow on the wall leap to an awful life before the light.

The next morning at breakfast Henry Glynn announced that he had to go to the city for three days. The sisters looked at him with surprise. He very seldom left home, and just now his practice had been neglected on account of Edward's death.

"How can you leave your patients now?" asked Mrs. Brigham wonderingly.

"I don't know how to, but there is no other way," replied Henry easily. "I have had a telegram from Dr. Mitford."

"Consultation?" inquired Mrs. Brigham.

"I have business," replied Henry.

Doctor Mitford was an old classmate of his who lived in a neighboring city and who occasionally called upon him in the case of a consultation.

After he had gone, Mrs. Brigham said to Caroline that, after all, Henry had not said that he was going to consult with Doctor

Mitford, and she thought it very strange.

"Everything is very strange," said Rebecca with a shudder.

"What do you mean?" inquired Caroline.

"Nothing," replied Rebecca.

Nobody entered the study that day, nor the next. The third day Henry was expected home, but he did not arrive and the last train from the city had come.

"I call it pretty queer work," said Mrs. Brigham. "The idea of a doctor leaving his patients at such a time as this, and the idea of a consultation lasting three days! There is no sense in it, and now he has not come. I don't understand it, for my part."

"I don't either," said Rebecca.

They were all in the south parlor. There was no light in the study; the door was ajar.

Presently Mrs. Brigham roseshe could not have told why; something seemed to impel hersome will outside her own. She went out of the room, again wrapping her rustling skirts round that she might pass noiselessly, and began pushing at the swollen door of the study.

"She has not got any lamp," said Rebecca in a shaking voice.

Caroline, who was writing letters, rose again, took the only remaining lamp in the room, and followed her sister. Rebecca had risen, but she stood trembling, not venturing to follow.

The doorbell rang, but the others did not hear it; it was on the south door on the other side of the house from the study. Rebecca, after hesitating until the bell rang the second time, went to the door; she remembered that the servant was out.

Caroline and her sister Emma entered the study. Caroline set the lamp on the table. They looked at the wall, and there were two shadows. The sisters stood clutching each other, staring at the awful things on the wall. Then Rebecca came in, staggering, with a telegram in her hand. "Here isa telegram," she gasped. "Henry isdead."

C.2.2 The Mass of Shadows

Source: <http://www.gutenberg.org/files/15143/15143-h/15143-h.htm#Mass>

This tale the sacristan of the church of St. Eulalie at Neuville d'Aumont told me, as we sat under the arbor of the White Horse, one fine summer evening, drinking a bottle of old wine to the health of the dead man, now very much at his ease, whom that very morning he had borne to the grave with full honors, beneath a pall powdered with smart silver tears.

"My poor father who is dead" (it is the sacristan who is speaking,) "was in his lifetime a grave-digger. He was of an agreeable disposition, the result, no doubt, of the calling he followed, for it has often been pointed out that people who work in cemeteries are of a jovial turn. Death has no terrors for them; they never give it a thought. I, for instance, monsieur, enter a cemetery at night as little perturbed as though it were the arbor of the White Horse. And if by chance I meet with a ghost, I don't disturb myself in the least about it, for I reflect that he may just as likely have business of his own to attend to as I. I know the habits of the dead,

and I know their character. Indeed, so far as that goes, I know things of which the priests themselves are ignorant. If I were to tell you all I have seen, you would be astounded. But a still tongue makes a wise head, and my father, who, all the same, delighted in spinning a yarn, did not disclose a twentieth part of what he knew. To make up for this he often repeated the same stories, and to my knowledge he told the story of Catherine Fontaine at least a hundred times.

"Catherine Fontaine was an old maid whom he well remembered having seen when he was a mere child. I should not be surprised if there were still, perhaps, three old fellows in the district who could remember having heard folks speak of her, for she was very well known and of excellent reputation, though poor enough. She lived at the corner of the Rue aux Nonnes, in the turret which is still to be seen there, and which formed part of an old half-ruined mansion looking on to the garden of the Ursuline nuns. On that turret can still be traced certain figures and half-obliterated inscriptions. The late cur of St. Eulalie, Monsieur Levasseur, asserted that there are the words in Latin, Love is stronger than death, 'which is to be understood,' so he would add, 'of divine love.'

"Catherine Fontaine lived by herself in this tiny apartment. She was a lace-maker. You know, of course, that the lace made in our part of the world was formerly held in high esteem. No one knew anything of her relatives or friends. It was reported that when she was eighteen years of age she had loved the young Chevalier d'Aumont-Clry, and had been secretly affianced to him. But decent folk didn't believe a word of it, and said it was nothing but a tale concocted because Catherine Fontaine's demeanor was that of a lady rather than that of a working woman, and because, moreover, she possessed beneath her white locks the remains of great beauty. Her expression was sorrowful, and on one finger she wore one of those rings fashioned by the goldsmith into the semblance of two tiny hands clasped together. In former days folks were accustomed to exchange such rings at their betrothal ceremony. I am sure you know the sort of thing I mean.

"Catherine Fontaine lived a saintly life. She spent a great deal of time in churches, and every morning, whatever might be the weather, she went to assist at the six o'clock Mass at St. Eulalie.

"Now one December night, whilst she was in her little chamber, she was awakened by the sound of bells, and nothing doubting that they were ringing for the first Mass, the pious woman dressed herself, and came downstairs and out into the street. The night was so obscure that not even the walls of the houses were visible, and not a ray of light shone from the murky sky. And such was the silence amid this black darkness, that there was not even the sound of a distant dog barking, and a feeling of aloofness from every living creature was perceptible. But Catherine Fontaine knew well every single stone she stepped on, and, as she could have found her way to the church with her eyes shut, she reached without difficulty the corner of the Rue aux Nonnes and the Rue de la Paroisse, where the timbered house stands with the tree of Jesse carved on one of its massive beams. When she reached this spot she perceived that the church doors were open, and that a great light was streaming out from the wax tapers. She resumed her journey, and when she had passed through the porch she found herself in the midst of a vast congregation which entirely filled the church. But she did not recognize any of the worshipers and was surprised to observe that all of these people were dressed in velvets and brocades, with feathers in their hats, and that they wore swords in the fashion of days gone by. Here were gentlemen who carried tall canes with gold knobs, and ladies with lace caps fastened with coronet-shaped combs. Chevaliers of the Order of St. Louis extended their hands to these ladies, who concealed behind their fans painted faces, of which only the powdered brow and the patch at the corner of the eye were visible! All of them proceeded to take their places without the slightest sound, and as they moved neither the sound of their footsteps on the pavement, nor the rustle of their garments could be heard. The lower places were filled with a crowd of young artisans in brown jackets, dimity breeches, and blue stockings, with their arms round the waists of pretty blushing girls who lowered their eyes. Near the holy water stoups peasant women, in scarlet petticoats and laced bodices, sat upon the ground as immovable as domestic animals, whilst young lads, standing up behind them, stared out from wide-open eyes and twirled their hats

round and round on their fingers, and all these sorrowful countenances seemed centred irremovably on one and the same thought, at once sweet and sorrowful. On her knees, in her accustomed place, Catherine Fontaine saw the priest advance toward the altar, preceded by two servers. She recognized neither priest nor clerks. The Mass began. It was a silent Mass, during which neither the sound of the moving lips nor the tinkle of the bell was audible. Catherine Fontaine felt that she was under the observation and the influence also of her mysterious neighbor, and when, scarcely turning her head, she stole a glance at him, she recognized the young Chevalier d'Aumont-Clry, who had once loved her, and who had been dead for five and forty years. She recognized him by a small mark which he had over the left ear, and above all by the shadow which his long black eyelashes cast upon his cheeks. He was dressed in his hunting clothes, scarlet with gold lace, the very clothes he wore that day when he met her in St. Leonard's Wood, begged of her a drink, and stole a kiss. He had preserved his youth and good looks. When he smiled, he still displayed magnificent teeth. Catherine said to him in an undertone:

"Monseigneur, you who were my friend, and to whom in days gone by I gave all that a girl holds most dear, may God keep you in His grace! O, that He would at length inspire me with regret for the sin I committed in yielding to you; for it is a fact that, though my hair is white and I approach my end, I have not yet repented of having loved you. But, dear dead friend and noble seigneur, tell me, who are these folk, habited after the antique fashion, who are here assisting at this silent Mass?"

"The Chevalier d'Aumont-Clry replied in a voice feebler than a breath, but none the less crystal clear:

"Catherine, these men and women are souls from purgatory who have grieved God by sinning as we ourselves sinned through love of the creature, but who are not on that account cast off by God, inasmuch as their sin, like ours, was not deliberate.

"Whilst separated from those whom they loved upon earth, they are purified in the cleansing fires of purgatory, they suffer the pangs of absence, which is for them the most cruel of tortures. They are so unhappy that an angel from heaven takes pity upon their love-torment. By the permission of the Most High, for one hour in the night, he reunites each year lover to loved in their parish church, where they are permitted to assist at the Mass of Shadows, hand clasped in hand. These are the facts. If it has been granted to me to see thee before thy death, Catherine, it is a boon which is bestowed by God's special permission."

"And Catherine Fontaine answered him:

"I would die gladly enough, dear, dead lord, if I might recover the beauty that was mine when I gave you to drink in the forest."

"Whilst they thus conversed under their breath, a very old canon was taking the collection and proffering to the worshipers a great copper dish, wherein they let fall, each in his turn, ancient coins which have long since ceased to pass current: cus of six livres, florins, ducats and ducatoons, jacobuses and rose-nobles, and the pieces fell silently into the dish. When at length it was placed before the Chevalier, he dropped into it a louis which made no more sound than had the other pieces of gold and silver.

"Then the old canon stopped before Catherine Fontaine, who fumbled in her pocket without being able to find a farthing. Then, being unwilling to allow the dish to pass without an offering from herself, she slipped from her finger the ring which the Chevalier had given her the day before his death, and cast it into the copper bowl. As the golden ring fell, a sound like the heavy clang of a bell rang out, and on the stroke of this reverberation the Chevalier, the canon, the celebrant, the servers, the ladies and their cavaliers, the whole assembly vanished utterly; the candles guttered out, and Catherine Fontaine was left alone in the darkness."

Having concluded his narrative after this fashion, the sacristan drank a long draught of wine, remained pensive for a moment, and then resumed his talk in these words:

"I have told you this tale exactly as my father has told it to me over and over again, and I believe that it is authentic, because it agrees in all respects with what I have observed of the manners and customs peculiar to those who have passed away. I have associated a good deal with the dead ever since my childhood, and I know that they are accustomed to return to what they have

loved.

"It is on this account that the miserly dead wander at night in the neighborhood of the treasures they conceal during their life time. They keep a strict watch over their gold; but the trouble they give themselves, far from being of service to them, turns to their disadvantage; and it is not a rare thing at all to come upon money buried in the ground on digging in a place haunted by a ghost. In the same way deceased husbands come by night to harass their wives who have made a second matrimonial venture, and I could easily name several who have kept a better watch over their wives since death than they ever did while living.

"That sort of thing is blameworthy, for in all fairness the dead have no business to stir up jealousies. Still I do but tell you what I have observed myself. It is a matter to take into account if one marries a widow. Besides, the tale I have told you is vouchsafed for in the manner following:

"The morning after that extraordinary night Catherine Fontaine was discovered dead in her chamber. And the beadle attached to St. Eulalie found in the copper bowl used for the collection a gold ring with two clasped hands. Besides, I'm not the kind of man to make jokes. Suppose we order another bottle of wine?..."

C.2.3 A Ghost

Source: <http://www.gutenberg.org/files/15143/15143-h/15143-h.htm#Ghosts>

We were speaking of sequestration, alluding to a recent lawsuit. It was at the close of a friendly evening in a very old mansion in the Rue de Grenelle, and each of the guests had a story to tell, which he assured us was true.

Then the old Marquis de la Tour-Samuel, eighty-two years of age, rose and came forward to lean on the mantelpiece. He told the following story in his slightly quavering voice.

"I, also, have witnessed a strange thing so strange that it has been the nightmare of my life. It happened fifty-six years ago, and yet there is not a month when I do not see it again in my dreams. From that day I have borne a mark, a stamp of fear, do you understand?

"Yes, for ten minutes I was a prey to terror, in such a way that ever since a constant dread has remained in my soul. Unexpected sounds chill me to the heart; objects which I can ill distinguish in the evening shadows make me long to flee. I am afraid at night.

"No! I would not have owned such a thing before reaching my present age. But now I may tell everything. One may fear imaginary dangers at eighty-two years old. But before actual danger I have never turned back, mesdames.

"That affair so upset my mind, filled me with such a deep, mysterious unrest that I never could tell it. I kept it in that inmost part, that corner where we conceal our sad, our shameful secrets, all the weaknesses of our life which cannot be confessed.

"I will tell you that strange happening just as it took place, with no attempt to explain it. Unless I went mad for one short hour it must be explainable, though. Yet I was not mad, and I will prove it to you. Imagine what you will. Here are the simple facts:

"It was in 1827, in July. I was quartered with my regiment in Rouen.

"One day, as I was strolling on the quay, I came across a man I believed I recognized, though I could not place him with certainty. I instinctively went more slowly, ready to pause. The stranger saw my impulse, looked at me, and fell into my arms.

"It was a friend of my younger days, of whom I had been very fond. He seemed to have become half a century older in the five years since I had seen him. His hair was white, and he stooped in his walk, as if he were exhausted. He understood my amazement and told me the story of his life.

"A terrible event had broken him down. He had fallen madly in love with a young girl and married her in a kind of dreamlike ecstasy. After a year of unalloyed bliss and unexhausted passion, she had died suddenly of heart disease, no doubt killed by love itself.

"He had left the country on the very day of her funeral, and had come to live in his hotel at Rouen. He remained there, solitary and desperate, grief slowly mining him, so wretched that he constantly thought of suicide.

"'As I thus came across you again,' he said, 'I shall ask a great favor of you. I want you to go to my chateau and get some papers I urgently need. They are in the writing-desk of my room, of our room. I cannot send a servant or a lawyer, as the errand must be kept private. I want absolute silence.

"'I shall give you the key of the room, which I locked carefully myself before leaving, and the key to the writing-desk. I shall also give you a note for the gardener, who will let you in.

"'Come to breakfast with me to-morrow, and we'll talk the matter over.'

"I promised to render him that slight service. It would mean but a pleasant excursion for me, his home not being more than twenty-five miles from Rouen. I could go there in an hour on horseback.

"At ten o'clock the next day I was with him. We breakfasted alone together, yet he did not utter more than twenty words. He asked me to excuse him. The thought that I was going to visit the room where his happiness lay shattered, upset him, he said. Indeed, he seemed perturbed, worried, as if some mysterious struggle were taking place in his soul.

"At last he explained exactly what I was to do. It was very simple. I was to take two packages of letters and some papers, locked in the first drawer at the right of the desk of which I had the key. He added:

"'I need not ask you not to glance at them.'

"I was almost hurt by his words, and told him so, rather sharply. He stammered:

"'Forgive me. I suffer so much!'

"And tears came to his eyes.

"I left about one o'clock to accomplish my errand.

"The day was radiant, and I rushed through the meadows, listening to the song of the larks, and the rhythmical beat of my sword on my riding-boots.

"Then I entered the forest, and I set my horse to walking. Branches of the trees softly caressed my face, and now and then I would catch a leaf between my teeth and bite it with avidity, full of the joy of life, such as fills you without reason, with a tumultuous happiness almost indefinable, a kind of magical strength.

"As I neared the house I took out the letter for the gardener, and noted with surprise that it was sealed. I was so amazed and so annoyed that I almost turned back without fulfilling my mission. Then I thought that I should thus display over-sensitiveness and bad taste. My friend might have sealed it unconsciously, worried as he was.

"The manor looked as though it had been deserted the last twenty years. The gate, wide-open and rotten, held, one wondered how. Grass filled the paths; you could not tell the flower-beds from the lawn.

"At the noise I made kicking a shutter, an old man came out from a side-door and was apparently amazed to see me there. I

dismounted from my horse and gave him the letter. He read it once or twice, turned it over, looked at me with suspicion, and asked:

"'Well, what do you want?'

"I answered sharply:

"'You must know it as you have read your master's orders. I want to get in the house.'

"He appeared overwhelmed. He said:

"'So you are going in in his room?'

"I was getting impatient.

"'Parbleu! Do you intend to question me, by chance?'

"He stammered:

"'Nomonsieuronlyit has not been opened since since the death. If you will wait five minutes, I will go in to see whether'

"I interrupted angrily:

"'See here, are you joking? You can't go in that room, as I have the key!'

"He no longer knew what to say.

"'Then, monsieur, I will show you the way.'

"'Show me the stairs and leave me alone. I can find it without your help.'

"'But stillmonsieur'

"Then I lost my temper.

"'Now be quiet! Else you'll be sorry!'

"I roughly pushed him aside and went into the house.

"I first went through the kitchen, then crossed two small rooms occupied by the man and his wife. From there I stepped into a large hall. I went up the stairs, and I recognized the door my friend had described to me.

"I opened it with ease and went in.

"The room was so dark that at first I could not distinguish anything. I paused, arrested by that moldy and stale odor peculiar to deserted and condemned rooms, of dead rooms. Then gradually my eyes grew accustomed to the gloom, and I saw rather clearly a great room in disorder, a bed without sheets having still its mattresses and pillows, one of which bore the deep print of an elbow or a head, as if someone had just been resting on it.

"The chairs seemed all in confusion. I noticed that a door, probably that of a closet, had remained ajar.

"I first went to the window and opened it to get some light, but the hinges of the outside shutters were so rusted that I could not loosen them.

"I even tried to break them with my sword, but did not succeed. As those fruitless attempts irritated me, and as my eyes were by now adjusted to the dim light, I gave up hope of getting more light and went toward the writing-desk.

"I sat down in an arm-chair, folded back the top, and opened the drawer. It was full to the edge. I needed but three packages, which I knew how to distinguish, and I started looking for them.

"I was straining my eyes to decipher the inscriptions, when I thought I heard, or rather felt a rustle behind me. I took no notice, thinking a draft had lifted some curtain. But a minute later, another movement, almost indistinct, sent a disagreeable little shiver over my skin. It was so ridiculous to be moved thus even so slightly, that I would not turn round, being ashamed. I had just discovered the second package I needed, and was on the point of reaching for the third, when a great and sorrowful sigh, close to my shoulder,

made me give a mad leap two yards away. In my spring I had turned round, my hand on the hilt of my sword, and surely had I not felt that, I should have fled like a coward.

"A tall woman, dressed in white, was facing me, standing behind the chair in which I had sat a second before.

"Such a shudder ran through me that I almost fell back! Oh, no one who has not felt them can understand those gruesome and ridiculous terrors! The soul melts; your heart seems to stop; your whole body becomes limp as a sponge, and your innermost parts seem collapsing.

"I do not believe in ghosts; and yet I broke down before the hideous fear of the dead; and I suffered, oh, I suffered more in a few minutes, in the irresistible anguish of supernatural dread, than I have suffered in all the rest of my life!

"If she had not spoken, I might have died. But she did speak; she spoke in a soft and plaintive voice which set my nerves vibrating. I could not say that I regained my self-control. No, I was past knowing what I did; but the kind of pride I have in me, as well as a military pride, helped me to maintain, almost in spite of myself, an honorable countenance. I was making a pose, a pose for myself, and for her, for her, whatever she was, woman, or phantom. I realized this later, for at the time of the apparition, I could think of nothing. I was afraid.

"She said:

"'Oh, you can be of great help to me, monsieur!'

"I tried to answer, but I was unable to utter one word. A vague sound came from my throat.

"She continued:

"'Will you? You can save me, cure me. I suffer terribly. I always suffer. I suffer, oh, I suffer!'

"And she sat down gently in my chair. She looked at me.

"'Will you?'

"I nodded my head, being still paralyzed.

"Then she handed me a woman's comb of tortoise-shell, and murmured:

"'Comb my hair! Oh, comb my hair! That will cure me. Look at my head how I suffer! And my hair how it hurts!'

"Her loose hair, very long, very black, it seemed to me, hung over the back of the chair, touching the floor.

"Why did I do it? Why did I, shivering, accept that comb, and why did I take between my hands her long hair, which left on my skin a ghastly impression of cold, as if I had handled serpents? I do not know.

"That feeling still clings about my fingers, and I shiver when I recall it.

"I combed her, I handled, I know not how, that hair of ice. I bound and unbound it; I plaited it as one plaits a horse's mane. She sighed, bent her head, seemed happy.

"Suddenly she said, 'Thank you!' tore the comb from my hands, and fled through the door which I had noticed was half opened.

"Left alone, I had for a few seconds the hazy feeling one feels in waking up from a nightmare. Then I recovered myself. I ran to the window and broke the shutters by my furious assault.

"A stream of light poured in. I rushed to the door through which that being had gone. I found it locked and immovable.

"Then a fever of flight seized on me, a panic, the true panic of battle. I quickly grasped the three packages of letters from the open desk; I crossed the room running, I took the steps of the stairway four at a time. I found myself outside, I don't know how, and seeing my horse close by, I mounted in one leap and left at a full gallop.

"I didn't stop till I reached Rouen and drew up in front of my house. Having thrown the reins to my orderly, I flew to my room and locked myself in to think.

"Then for an hour I asked myself whether I had not been the victim of an hallucination. Certainly I must have had one of those nervous shocks, one of those brain disorders such as give rise to miracles, to which the supernatural owes its strength.

"And I had almost concluded that it was a vision, an illusion of my senses, when I came near to the window. My eyes by chance looked down. My tunic was covered with hairs, long woman's hairs which had entangled themselves around the buttons!

"I took them off one by one and threw them out of the window with trembling fingers.

"I then called my orderly. I felt too perturbed, too moved, to go and see my friend on that day. Besides, I needed to think over what I should tell him.

"I had his letters delivered to him. He gave a receipt to the soldier. He inquired after me and was told that I was not well. I had had a sunstroke, or something. He seemed distressed.

"I went to see him the next day, early in the morning, bent on telling him the truth. He had gone out the evening before and had not come back.

"I returned the same day, but he had not been seen. I waited a week. He did not come back. I notified the police. They searched for him everywhere, but no one could find any trace of his passing or of his retreat.

"A careful search was made in the deserted manor. No suspicious clue was discovered.

"There was no sign that a woman had been concealed there.

"The inquest gave no result, and so the search went no further.

"And in fifty-six years I have learned nothing more. I never found out the truth."

C.2.4 What Was It?

Source: <http://www.gutenberg.org/files/15143/15143-h/15143-h.htm#What>

It is, I confess, with considerable diffidence, that I approach the strange narrative which I am about to relate. The events which I purpose detailing are of so extraordinary a character that I am quite prepared to meet with an unusual amount of incredulity and scorn. I accept all such beforehand. I have, I trust, the literary courage to face unbelief. I have, after mature consideration resolved to narrate, in as simple and straightforward a manner as I can compass, some facts that passed under my observation, in the month of July last, and which, in the annals of the mysteries of physical science, are wholly unparalleled.

I live at No. Twenty-sixth Street, in New York. The house is in some respects a curious one. It has enjoyed for the last two years the reputation of being haunted. It is a large and stately residence, surrounded by what was once a garden, but which is now only a green enclosure used for bleaching clothes. The dry basin of what has been a fountain, and a few fruit trees ragged and unpruned, indicate that this spot in past days was a pleasant, shady retreat, filled with fruits and flowers and the sweet murmur of waters.

The house is very spacious. A hall of noble size leads to a large spiral staircase winding through its center, while the various apartments are of imposing dimensions. It was built some fifteen or twenty years since by Mr. A, the well-known New York merchant, who five years ago threw the commercial world into convulsions by a stupendous bank fraud. Mr. A, as everyone knows, escaped to Europe, and died not long after, of a broken heart. Almost immediately after the news of his decease reached this country and was

verified, the report spread in Twenty-sixth Street that No. was haunted. Legal measures had dispossessed the widow of its former owner, and it was inhabited merely by a caretaker and his wife, placed there by the house agent into whose hands it had passed for the purposes of renting or sale. These people declared that they were troubled with unnatural noises. Doors were opened without any visible agency. The remnants of furniture scattered through the various rooms were, during the night, piled one upon the other by unknown hands. Invisible feet passed up and down the stairs in broad daylight, accompanied by the rustle of unseen silk dresses, and the gliding of viewless hands along the massive balusters. The caretaker and his wife declared they would live there no longer. The house agent laughed, dismissed them, and put others in their place. The noises and supernatural manifestations continued. The neighborhood caught up the story, and the house remained untenanted for three years. Several persons negotiated for it; but, somehow, always before the bargain was closed they heard the unpleasant rumors and declined to treat any further.

It was in this state of things that my landlady, who at that time kept a boarding-house in Bleeker Street, and who wished to move further up town, conceived the bold idea of renting No. Twenty-sixth Street. Happening to have in her house rather a plucky and philosophical set of boarders, she laid her scheme before us, stating candidly everything she had heard respecting the ghostly qualities of the establishment to which she wished to remove us. With the exception of two timid persons, a sea-captain and a returned Californian, who immediately gave notice that they would leave, all of Mrs. Moffat's guests declared that they would accompany her in her chivalric incursion into the abode of spirits.

Our removal was effected in the month of May, and we were charmed with our new residence. The portion of Twenty-sixth Street where our house is situated, between Seventh and Eighth Avenues, is one of the pleasantest localities in New York. The gardens back of the houses, running down nearly to the Hudson, form, in the summer time, a perfect avenue of verdure. The air is pure and invigorating, sweeping, as it does, straight across the river from the Weehawken heights, and even the ragged garden which surrounded the house, although displaying on washing days rather too much clothesline, still gave us a piece of greensward to look at, and a cool retreat in the summer evenings, where we smoked our cigars in the dusk, and watched the fireflies flashing their dark lanterns in the long grass.

Of course we had no sooner established ourselves at No. than we began to expect ghosts. We absolutely awaited their advent with eagerness. Our dinner conversation was supernatural. One of the boarders, who had purchased Mrs. Crowe's *Night Side of Nature* for his own private delectation, was regarded as a public enemy by the entire household for not having bought twenty copies. The man led a life of supreme wretchedness while he was reading this volume. A system of espionage was established, of which he was the victim. If he incautiously laid the book down for an instant and left the room, it was immediately seized and read aloud in secret places to a select few. I found myself a person of immense importance, it having leaked out that I was tolerably well versed in the history of supernaturalism, and had once written a story the foundation of which was a ghost. If a table or a wainscot panel happened to warp when we were assembled in the large drawing-room, there was an instant silence, and everyone was prepared for an immediate clanking of chains and a spectral form.

After a month of psychological excitement, it was with the utmost dissatisfaction that we were forced to acknowledge that nothing in the remotest degree approaching the supernatural had manifested itself. Once the black butler asseverated that his candle had been blown out by some invisible agency while he was undressing himself for the night; but as I had more than once discovered this colored gentleman in a condition when one candle must have appeared to him like two, thought it possible that, by going a step further in his potations, he might have reversed this phenomenon, and seen no candle at all where he ought to have beheld one.

Things were in this state when an accident took place so awful and inexplicable in its character that my reason fairly reels at the bare memory of the occurrence. It was the tenth of July. After dinner was over I repaired, with my friend Dr. Hammond, to

the garden to smoke my evening pipe. Independent of certain mental sympathies which existed between the Doctor and myself, we were linked together by a vice. We both smoked opium. We knew each other's secret, and respected it. We enjoyed together that wonderful expansion of thought, that marvelous intensifying of the perceptive faculties, that boundless feeling of existence when we seem to have points of contact with the whole universe, in short, that unimaginable spiritual bliss, which I would not surrender for a throne, and which I hope you, reader, will never taste.

Those hours of opium happiness which the Doctor and I spent together in secret were regulated with a scientific accuracy. We did not blindly smoke the drug of paradise, and leave our dreams to chance. While smoking, we carefully steered our conversation through the brightest and calmest channels of thought. We talked of the East, and endeavored to recall the magical panorama of its glowing scenery. We criticized the most sensuous poets, those who painted life ruddy with health, brimming with passion, happy in the possession of youth and strength and beauty. If we talked of Shakespeare's *Tempest*, we lingered over Ariel, and avoided Caliban. Like the Guebers, we turned our faces to the East, and saw only the sunny side of the world.

This skillful coloring of our train of thought produced in our subsequent visions a corresponding tone. The splendors of Arabian fairyland dyed our dreams. We paced the narrow strip of grass with the tread and port of kings. The song of the Rana arborea, while he clung to the bark of the ragged plum-tree, sounded like the strains of divine musicians. Houses, walls, and streets melted like rain clouds, and vistas of unimaginable glory stretched away before us. It was a rapturous companionship. We enjoyed the vast delight more perfectly because, even in our most ecstatic moments, we were conscious of each other's presence. Our pleasures, while individual, were still twin, vibrating and moving in musical accord.

On the evening in question, the tenth of July, the Doctor and myself drifted into an unusually metaphysical mood. We lit our large meerschaums, filled with fine Turkish tobacco, in the core of which burned a little black nut of opium, that, like the nut in the fairy tale, held within its narrow limits wonders beyond the reach of kings; we paced to and fro, conversing. A strange perversity dominated the currents of our thought. They would not flow through the sun-lit channels into which we strove to divert them. For some unaccountable reason, they constantly diverged into dark and lonesome beds, where a continual gloom brooded. It was in vain that, after our old fashion, we flung ourselves on the shores of the East, and talked of its gay bazaars, of the splendors of the time of Haroun, of harems and golden palaces. Black afreets continually arose from the depths of our talk, and expanded, like the one the fisherman released from the copper vessel, until they blotted everything bright from our vision. Insensibly, we yielded to the occult force that swayed us, and indulged in gloomy speculation. We had talked some time upon the proneness of the human mind to mysticism, and the almost universal love of the terrible, when Hammond suddenly said to me. "What do you consider to be the greatest element of terror?"

The question puzzled me. That many things were terrible, I knew. Stumbling over a corpse in the dark; beholding, as I once did, a woman floating down a deep and rapid river, with wildly lifted arms, and awful, upturned face, uttering, as she drifted, shrieks that rent one's heart while we, spectators, stood frozen at a window which overhung the river at a height of sixty feet, unable to make the slightest effort to save her, but dumbly watching her last supreme agony and her disappearance. A shattered wreck, with no life visible, encountered floating listlessly on the ocean, is a terrible object, for it suggests a huge terror, the proportions of which are veiled. But it now struck me, for the first time, that there must be one great and ruling embodiment of fear, a King of Terrors, to which all others must succumb. What might it be? To what train of circumstances would it owe its existence?

"I confess, Hammond," I replied to my friend, "I never considered the subject before. That there must be one Something more terrible than any other thing, I feel. I cannot attempt, however, even the most vague definition."

"I am somewhat like you, Harry," he answered. "I feel my capacity to experience a terror greater than anything yet conceived by

the human mind; something combining in fearful and unnatural amalgamation hitherto supposed incompatible elements. The calling of the voices in Brockden Brown's novel of *Wieland* is awful; so is the picture of the Dweller of the Threshold, in Bulwer's *Zanoni*; but," he added, shaking his head gloomily, "there is something more horrible still than those."

"Look here, Hammond," I rejoined, "let us drop this kind of talk, for Heaven's sake! We shall suffer for it, depend on it."

"I don't know what's the matter with me to-night," he replied, "but my brain is running upon all sorts of weird and awful thoughts. I feel as if I could write a story like Hoffman, to-night, if I were only master of a literary style."

"Well, if we are going to be Hoffmanesque in our talk, I'm off to bed. Opium and nightmares should never be brought together. How sultry it is! Good-night, Hammond."

"Good-night, Harry. Pleasant dreams to you."

"To you, gloomy wretch, afreets, ghouls, and enchanters."

We parted, and each sought his respective chamber. I undressed quickly and got into bed, taking with me, according to my usual custom, a book, over which I generally read myself to sleep. I opened the volume as soon as I had laid my head upon the pillow, and instantly flung it to the other side of the room. It was Goudon's *History of Monsters*, a curious French work, which I had lately imported from Paris, but which, in the state of mind I had then reached, was anything but an agreeable companion. I resolved to go to sleep at once; so, turning down my gas until nothing but a little blue point of light glimmered on the top of the tube, I composed myself to rest.

The room was in total darkness. The atom of gas that still remained alight did not illuminate a distance of three inches round the burner. I desperately drew my arm across my eyes, as if to shut out even the darkness, and tried to think of nothing. It was in vain. The confounded themes touched on by Hammond in the garden kept obtruding themselves on my brain. I battled against them. I erected ramparts of would-be blackness of intellect to keep them out. They still crowded upon me. While I was lying still as a corpse, hoping that by a perfect physical inaction I should hasten mental repose, an awful incident occurred. A Something dropped, as it seemed, from the ceiling, plumb upon my chest, and the next instant I felt two bony hands encircling my throat, endeavoring to choke me.

I am no coward, and am possessed of considerable physical strength. The suddenness of the attack, instead of stunning me, strung every nerve to its highest tension. My body acted from instinct, before my brain had time to realize the terrors of my position. In an instant I wound two muscular arms around the creature, and squeezed it, with all the strength of despair, against my chest. In a few seconds the bony hands that had fastened on my throat loosened their hold, and I was free to breathe once more. Then commenced a struggle of awful intensity. Immersed in the most profound darkness, totally ignorant of the nature of the Thing by which I was so suddenly attacked, finding my grasp slipping every moment, by reason, it seemed to me, of the entire nakedness of my assailant, bitten with sharp teeth in the shoulder, neck, and chest, having every moment to protect my throat against a pair of sinewy, agile hands, which my utmost efforts could not confine, these were a combination of circumstances to combat which required all the strength, skill, and courage that I possessed.

At last, after a silent, deadly, exhausting struggle, I got my assailant under by a series of incredible efforts of strength. Once pinned, with my knee on what I made out to be its chest, I knew that I was victor. I rested for a moment to breathe. I heard the creature beneath me panting in the darkness, and felt the violent throbbing of a heart. It was apparently as exhausted as I was; that was one comfort. At this moment I remembered that I usually placed under my pillow, before going to bed, a large yellow silk pocket handkerchief. I felt for it instantly; it was there. In a few seconds more I had, after a fashion, pinioned the creature's arms.

I now felt tolerably secure. There was nothing more to be done but to turn on the gas, and, having first seen what my midnight

assailant was like, arouse the household. I will confess to being actuated by a certain pride in not giving the alarm before; I wished to make the capture alone and unaided.

Never losing my hold for an instant, I slipped from the bed to the floor, dragging my captive with me. I had but a few steps to make to reach the gas-burner; these I made with the greatest caution, holding the creature in a grip like a vice. At last I got within arm's length of the tiny speck of blue light which told me where the gas-burner lay. Quick as lightning I released my grasp with one hand and let on the full flood of light. Then I turned to look at my captive.

I cannot even attempt to give any definition of my sensations the instant after I turned on the gas. I suppose I must have shrieked with terror, for in less than a minute afterward my room was crowded with the inmates of the house. I shudder now as I think of that awful moment. I saw nothing! Yes; I had one arm firmly clasped round a breathing, panting, corporeal shape, my other hand gripped with all its strength a throat as warm, as apparently fleshy, as my own; and yet, with this living substance in my grasp, with its body pressed against my own, and all in the bright glare of a large jet of gas, I absolutely beheld nothing! Not even an outline, a vapor!

I do not, even at this hour, realize the situation in which I found myself. I cannot recall the astounding incident thoroughly. Imagination in vain tries to compass the awful paradox.

It breathed. I felt its warm breath upon my cheek. It struggled fiercely. It had hands. They clutched me. Its skin was smooth, like my own. There it lay, pressed close up against me, solid as stone, and yet utterly invisible!

I wonder that I did not faint or go mad on the instant. Some wonderful instinct must have sustained me; for, absolutely, in place of loosening my hold on the terrible Enigma, I seemed to gain an additional strength in my moment of horror, and tightened my grasp with such wonderful force that I felt the creature shivering with agony.

Just then Hammond entered my room at the head of the household. As soon as he beheld my face which, I suppose, must have been an awful sight to look at he hastened forward, crying, "Great heaven, Harry! what has happened?"

"Hammond! Hammond!" I cried, "come here. O, this is awful! I have been attacked in bed by something or other, which I have hold of; but I can't see it, I can't see it!"

Hammond, doubtless struck by the unfeigned horror expressed in my countenance, made one or two steps forward with an anxious yet puzzled expression. A very audible titter burst from the remainder of my visitors. This suppressed laughter made me furious. To laugh at a human being in my position! It was the worst species of cruelty. Now, I can understand why the appearance of a man struggling violently, as it would seem, with an airy nothing, and calling for assistance against a vision, should have appeared ludicrous. Then, so great was my rage against the mocking crowd that had I the power I would have stricken them dead where they stood.

"Hammond! Hammond!" I cried again, despairingly, "for God's sake come to me. I can hold the thing but a short while longer. It is overpowering me. Help me! Help me!"

"Harry," whispered Hammond, approaching me, "you have been smoking too much opium."

"I swear to you, Hammond, that this is no vision," I answered, in the same low tone. "Don't you see how it shakes my whole frame with its struggles? If you don't believe me, convince yourself. Feel it, touch it."

Hammond advanced and laid his hand in the spot I indicated. A wild cry of horror burst from him. He had felt it!

In a moment he had discovered somewhere in my room a long piece of cord, and was the next instant winding it and knotting it about the body of the unseen being that I clasped in my arms.

"Harry," he said, in a hoarse, agitated voice, for, though he preserved his presence of mind, he was deeply moved, "Harry, it's all safe now. You may let go, old fellow, if you're tired. The Thing can't move."

I was utterly exhausted, and I gladly loosed my hold.

Hammond stood holding the ends of the cord that bound the Invisible, twisted round his hand, while before him, self-supporting as it were, he beheld a rope laced and interlaced, and stretching tightly around a vacant space. I never saw a man look so thoroughly stricken with awe. Nevertheless his face expressed all the courage and determination which I knew him to possess. His lips, although white, were set firmly, and one could perceive at a glance that, although stricken with fear, he was not daunted.

The confusion that ensued among the guests of the house who were witnesses of this extraordinary scene between Hammond and myself, who beheld the pantomime of binding this struggling Something, who beheld me almost sinking from physical exhaustion when my task of jailer was over, the confusion and terror that took possession of the bystanders, when they saw all this, was beyond description. The weaker ones fled from the apartment. The few who remained clustered near the door and could not be induced to approach Hammond and his Charge. Still incredulity broke out through their terror. They had not the courage to satisfy themselves, and yet they doubted. It was in vain that I begged of some of the men to come near and convince themselves by touch of the existence in that room of a living being which was invisible. They were incredulous, but did not dare to undeceive themselves. How could a solid, living, breathing body be invisible, they asked. My reply was this. I gave a sign to Hammond, and both of us conquering our fearful repugnance to touch the invisible creature lifted it from the ground, manacled as it was, and took it to my bed. Its weight was about that of a boy of fourteen.

"Now my friends," I said, as Hammond and myself held the creature suspended over the bed, "I can give you self-evident proof that here is a solid, ponderable body, which, nevertheless, you cannot see. Be good enough to watch the surface of the bed attentively."

I was astonished at my own courage in treating this strange event so calmly; but I had recovered from my first terror, and felt a sort of scientific pride in the affair, which dominated every other feeling.

The eyes of the bystanders were immediately fixed on my bed. At a given signal Hammond and I let the creature fall. There was a dull sound of a heavy body alighting on a soft mass. The timbers of the bed creaked. A deep impression marked itself distinctly on the pillow, and on the bed itself. The crowd who witnessed this gave a low cry, and rushed from the room. Hammond and I were left alone with our Mystery.

We remained silent for some time, listening to the low, irregular breathing of the creature on the bed, and watching the rustle of the bedclothes as it impotently struggled to free itself from confinement. Then Hammond spoke.

"Harry, this is awful."

"Ay, awful."

"But not unaccountable."

"Not unaccountable! What do you mean? Such a thing has never occurred since the birth of the world. I know not what to think, Hammond. God grant that I am not mad, and that this is not an insane fantasy!"

"Let us reason a little, Harry. Here is a solid body which we touch, but which we cannot see. The fact is so unusual that it strikes us with terror. Is there no parallel, though, for such a phenomenon? Take a piece of pure glass. It is tangible and transparent. A certain chemical coarseness is all that prevents its being so entirely transparent as to be totally invisible. It is not theoretically impossible, mind you, to make a glass which shall not reflect a single ray of light, a glass so pure and homogeneous in its atoms that the rays from the sun will pass through it as they do through the air, refracted but not reflected. We do not see the air, and yet we feel it."

"That's all very well, Hammond, but these are inanimate substances. Glass does not breathe, air does not breathe. This thing has a heart that palpitates, a will that moves it, lungs that play, and inspire and respire."

"You forget the phenomena of which we have so often heard of late," answered the Doctor, gravely. "At the meetings called 'spirit circles,' invisible hands have been thrust into the hands of those persons round the table, warm, fleshly hands that seemed to pulsate with mortal life."

"What? Do you think, then, that this thing is"

"I don't know what it is," was the solemn reply; "but please the gods I will, with your assistance, thoroughly investigate it."

We watched together, smoking many pipes, all night long, by the bedside of the unearthly being that tossed and panted until it was apparently wearied out. Then we learned by the low, regular breathing that it slept.

The next morning the house was all astir. The boarders congregated on the landing outside my room, and Hammond and myself were lions. We had to answer a thousand questions as to the state of our extraordinary prisoner, for as yet not one person in the house except ourselves could be induced to set foot in the apartment.

The creature was awake. This was evidenced by the convulsive manner in which the bedclothes were moved in its efforts to escape. There was something truly terrible in beholding, as it were, those second-hand indications of the terrible writhings and agonized struggles for liberty which themselves were invisible.

Hammond and myself had racked our brains during the long night to discover some means by which we might realize the shape and general appearance of the Enigma. As well as we could make out by passing our hands over the creature's form, its outlines and lineaments were human. There was a mouth; a round, smooth head without hair; a nose, which, however, was little elevated above the cheeks; and its hands and feet felt like those of a boy. At first we thought of placing the being on a smooth surface and tracing its outlines with chalk, as shoemakers trace the outline of the foot. This plan was given up as being of no value. Such an outline would give not the slightest idea of its conformation.

A happy thought struck me. We would take a cast of it in plaster of Paris. This would give us the solid figure, and satisfy all our wishes. But how to do it? The movements of the creature would disturb the setting of the plastic covering, and distort the mold. Another thought. Why not give it chloroform? It had respiratory organs, that was evident by its breathing. Once reduced to a state of insensibility, we could do with it what we would. Doctor X was sent for; and after the worthy physician had recovered from the first shock of amazement, he proceeded to administer the chloroform. In three minutes afterward we were enabled to remove the fetters from the creature's body, and a modeler was busily engaged in covering the invisible form with the moist clay. In five minutes more we had a mold, and before evening a rough facsimile of the Mystery. It was shaped like a man, distorted, uncouth, and horrible, but still a man. It was small, not over four feet and some inches in height, and its limbs revealed a muscular development that was unparalleled. Its face surpassed in hideousness anything I had ever seen. Gustav Dor, or Callot, or Tony Johannot, never conceived anything so horrible. There is a face in one of the latter's illustrations to *Un Voyage o il vous plaira*, which somewhat approaches the countenance of this creature, but does not equal it. It was the physiognomy of what I should fancy a ghoul might be. It looked as if it was capable of feeding on human flesh.

Having satisfied our curiosity, and bound every one in the house to secrecy, it became a question what was to be done with our Enigma? It was impossible that we should keep such a horror in our house; it was equally impossible that such an awful being should be let loose upon the world. I confess that I would have gladly voted for the creature's destruction. But who would shoulder the responsibility? Who would undertake the execution of this horrible semblance of a human being? Day after day this question was deliberated gravely. The boarders all left the house. Mrs. Moffat was in despair, and threatened Hammond and myself with all sorts of legal penalties if we did not remove the Horror. Our answer was, "We will go if you like, but we decline taking this creature with us. Remove it yourself if you please. It appeared in your house. On you the responsibility rests." To this there was, of course, no

answer. Mrs. Moffat could not obtain for love or money a person who would even approach the Mystery.

The most singular part of the affair was that we were entirely ignorant of what the creature habitually fed on. Everything in the way of nutriment that we could think of was placed before it, but was never touched. It was awful to stand by, day after day, and see the clothes toss, and hear the hard breathing, and know that it was starving.

Ten, twelve days, a fortnight passed, and it still lived. The pulsations of the heart, however, were daily growing fainter, and had now nearly ceased. It was evident that the creature was dying for want of sustenance. While this terrible life-struggle was going on, I felt miserable. I could not sleep. Horrible as the creature was, it was pitiful to think of the pangs it was suffering.

At last it died. Hammond and I found it cold and stiff one morning in the bed. The heart had ceased to beat, the lungs to inspire. We hastened to bury it in the garden. It was a strange funeral, the dropping of that viewless corpse into the damp hole. The cast of its form I gave to Doctor X, who keeps it in his museum in Tenth Street.

As I am on the eve of a long journey from which I may not return, I have drawn up this narrative of an event the most singular that has ever come to my knowledge.

C.3 Literature

C.3.1 Romeo and Juliet

Source: <https://www.booksummary.net/romeo-and-juliet-william-shakespeare/>

The plot begins on a warm night in July when the Capuletis servants stroll through town looking for troubles.

They find what they were looking for when they encounter the servants from the rival family Montecchi. Soon a fight breaks out and it leads to a sword battle between the servants. The confrontation was so big that at the end even the rival families took part in it. The duke Scala settled the fight by threatening to punish the servants if a disorder like that ever occurred again.

Romeos parents, the Montecchis, were happy that the fight did not include their son. Despite that Romeos behavior has been strange lately. He has been noticed walking through the forest at night, he avoids his friends and during the day he stays locked up in his bedroom. The concerned parents try to get the truth out of his cousin Benvolio.

Benvolio was not just Romeos cousin but also his best friend. The truth, alongside the real reason for Romeos strange behavior, comes out when Benvolio finds out that Romeo is unhappily in love with Rosaline. To cheer him up again, Benvolio suggests that they sneak into the Capuletis ball while convincing him that the cure for his disease is there. Even though he found that suggestion weird, he accepts.

In the meantime, a lot happens in the Capuleti family and the excitement increases by the moment. It is not only the ball that keeps them occupied, but also the arrival of the count Paris who is coming to ask for the hand of their daughter Juliet. Juliets father considers her too young to get married but does not hide his satisfaction with the fact that the count Paris is interested in his

daughter. Juliet is a thirteen years old girl with no experience in love but she gives her word to her father that she will try to seduce the count at the ball.

Juliet meets Romeo the same night and because of that Rosaline, alongside count Paris, gets forgotten. The special moment between the newly in love does not last long because Tybalt, the nephew of miss Capuleti, recognizes the disguised Romeo and wants to confront him. The old Capuleti manages to calm them down but the thirst for revenge was still very much alive. Soon after the ball and the misfortunate event, Romeo finds himself under Juliets window and in that glorious moment they promise each other eternal love and decide to get married.

Friar Laurence, despite his concern about this sudden love, decides to marry them. He hopes that the marriage between Romeo and Juliet will bring peace and forever end the family war. Unfortunately, thing didnt go the way they were planned. The recently married Romeo encounters his friends Benvolio and Mercutio who got into an argument with Tybalt who came looking for revenge.

Tyblat challenges Romeo to a fight, but Romeo says no. He admits that he loves the Capuleti as much as he loves his family. The confession caused shock. Mercutio accepts the challenge and Tybalt stabs him. Rage and the feeling of guilty force Romeo to attack Tybalt and they fight to death. Romeo is the one who gets out alive but he has to leave town because he committed murder. Juliet was excited about her first weeding night when the maid came and told her what happened. Juliet was devastated and admits to her maid that Romeo is hiding at friar Laurence.

The old Capuleti cant believe that Tybalt is dead and he is even more upset over the suffering of his daughter. Because of that he decides to rush her weeding with the count Paris. Juliet makes her father mad by refusing to marry the count. After she didnt find any understanding or help from anyone, Juliet goes to friar Laurence and comes up with a deadly plan.

She was going to drink a potion that will stop her breathing for 42 hours convincing everyone that she is dead. In the meantime friar Laurence will send a messenger to get Romeo who was hiding in Mantova and the two of them will, hidden in the tomb, wait for Juliet to wake up. Afterwards Romeo will take her with him and when friar Laurence declares that they are married, he will be able to return to Verona. Juliet drinks the potion.

The next morning when the maid came to help the bride get ready for the weeding she was shocked to find Juliet dead. The whole Capuleti house grieves because of their dead Juliet. Friar Laurence sticks to the plan and sends the messenger but he does not manage to get to Mantova in time. Romeos servant Balthasar gets to him first and lets him know that Juliet is dead. Romeo, overwhelmed by the news, buys poison and heads towards Verona. He finds Paris mourning Juliet and another fight breaks out that costs Paris his life.

Parises last wish was to be buried next to Juliet and Romeo promised him to do that. When he saw Juliet lying Romeo thought that she was actually dead and drinks the poison. Friar Laurence runs towards the tomb but he came too late. Romeo and Paris were already dead. Juliet wakes up and when she sees that Romeo is dead, she takes a dagger and stabs herself. The tragic death united the two families and put an end to the hate between them.

C.3.2 Harry Potter and The Sorcerers Stone

Source: <https://www.booksummary.net/harry-potter-and-the-sorcerers-stone-j-k-rowling/>

Harry Potter and the Sorcerers Stone, by J. K. Rowling opens in a normal suburb in a town in England. The Dursleys are a normal, higher income family, living in a normal house with a normal baby boy, Dudley. Mrs. Dursley, Petunia, does have a sister, Lilly, who is a bit eccentric and married to James Potter, who is also questionable. But, as long as the Dursleys dont have to acknowledge her, everything is fine.

One fine day, Mr. Dursley, Vernon, is coming home from work when he notices a few odd people about. Since he is the kind of gentleman who ignores those who are different and/or of a lower income bracket, he hurries past them. Even seeing a cat apparently reading a map and overhearing a conversation in front of the bakery that seems to mention his sister and brother-in-law, dont make him pause. That night, the Dursleys learn of some strange reports in the news of owls and shooting stars, but , they shrug it off and go to bed.

Meanwhile, as the streetlights go out, an old man in a strange outfit is walking down the street toward their house. He is Albus Dumbledore, the head wizard of the Hogwarts Wizarding Academy. A cat comes up to him, changing into a woman as it does. She is Professor McGonagall, of the same academy. Dumbledore tells her that you-know-who, or Voldemort, as he keeps trying to persuade people to use, has killed Lily and James Potter, little Harrys parents. When he tried to kill Henry, something happened, and his power began to dim. Voldemort ran away afterward.

Voldemort is the villain of this book. A fallen, but powerful wizard, he is in search of ultimate power and immortality. Therefore, he wants the Sorcerers Stone. The fabled stone that the alchemists are always looking for in order to make the elixir for immortality. Voldemort is so feared that wizards dont say his name, for fear he will hear them. He is always referred to as you-know-who.

Dumbledore and McGonagall are talking about James and Lilies one-year-old son, Harry. Having spent the day observing the Dursleys, McGonagall is against leaving him with them. But, Dumbledore insists it is the safest place for him. He says that when Harry is old enough, the boys aunt and uncle will tell him about the prophecy. He wrote them a letter, explaining the whole thing.

After the giant, Hagrid, arrives from the sky, on a motorcycle carrying the baby, Dumbledore places him on the doorstep of the Dursleys with the letter of explanation. Then Hagrid wipes his eyes, mounts the motorcycle and leaves, and McGonagall blew her nose, then left, also. Finally, Dumbledore returned the lights to their streetlamps and left with a swirl of his cloak. All this time, baby Harry slept only to be awakened in the morning by his aunts screams when she opened the door for the milk delivery. But, all night, people toasted to him, The boy who lived.

Ten years later, Harry is not living the life of a hero. He is an eleven-year-old boy, who is regularly abused by his aunt, uncle and especially their disgusting son, Dudley. He is spoiled and cruel. Dudley has always pinched and hit Harry. Today is Dudleys birthday and he is angry because he received one less gift than the year before, and since the next door neighbor cant watch Harry, he will be going along on Dudleys birthday trip to the zoo.

While at the zoo Harry is scorned by his aunt and uncle, and by Dudley and his friend. While in the snake habitat, Harry is surprised to discover he can speak with the boa constrictor. Dudley comes up looking for mischief and suddenly the glass in front of the boas cage disappears. Although Harry swears he had nothing to do with it, he is sent to the cupboard he sleeps in under the stairs, without any food. There he stays until summer. He is only allowed out for school and chores.

When summer arrives, Harry spends most of his free time outside. His cousin and his friends are inside, and they like to torment

him. One day, a letter arrives addressed to him in The Cupboard under the Stairs. When his uncle sees it, he grabs it. He and his wife make the children leave so they can discuss the letter, then later, his uncle tells him to move into the small room next to Dudley's that held his toys.

The next day, another letter arrives, addressed to the smallest bedroom. Harry's uncle is enraged. No matter what steps he takes to keep the letter from Harry, it just keeps coming. Hundreds of letters arrive addressed to Harry. Finally, he takes the family to a remote island and bolts the door, sure that he has thwarted the letters. Then at midnight, just as it turns to Harry's birthday, they hear a banging at the door.

Hagrid, the giant and Keeper of the Keys at Hogwarts, breaks down the door. When the uncle tries to shoot him with a shotgun, he takes it and ties it into a knot. Hagrid is very upset to learn that his aunt and uncle haven't told Harry about Hogwarts and that he is a wizard. Even though his uncle tries to stop him, Hagrid gives Harry the letter of acceptance for Hogwarts School of Witchcraft and Wizardry. Then Hagrid tells Harry the truth about how his parents died, not in a car accident like his aunt and uncle told him. Against his uncle's protests, Hagrid takes Harry away.

The next morning, Harry is surprised to learn the night before was not a dream and he is still with the giant. Hagrid takes him towards London to shop for school supplies. Harry is concerned about costs, but Hagrid tells him that his parents left him with plenty of money. They go to the Leaky Cauldron pub and on through to a brick wall. Hagrid taps on it and the bricks part revealing Diagon Alley. The street is lined with odd shops and busy with shoppers. The first stop is Gringotts Bank run by goblins, where Harry sees his parents' vault full of gold and silver. Hagrid helps him remove enough for his needs, then Hagrid makes a stop at vault #713, where he picks up a dirty little package, tucks it in his pocket and tells Harry not to ask about it.

While Harry is being fitted for his school uniform, he meets a snobby boy about his age. As the boy is talking about Hogwarts and Quidditch, a kind of soccer played with four balls and on broomsticks in the air, Harry feels less and less prepared for life as a wizard. Then Hagrid takes him to be fitted for his wand. After trying a few, the one that works for Harry is made of holly and has a phoenix feather. The shop owner tells him the only other wand with a feather from that phoenix was owned by Voldemort. It is the same wand that gave Harry the scar on his forehead shaped like a lightning bolt.

Finally, the day comes to leave for school. Harry's ticket says his train leaves on track 9. Unsure how to find the track, he overhears a family speaking of Hogwarts. Harry asks the mother for help and she shows him how to go through the barrier between tracks nine and ten. It is a family of red-haired children and three of the boys are headed to Hogwarts. Twin boys, Fred and George Weasley are returning, and Ron Weasley is starting like Harry is. Ron and Harry become fast friends as Ron explains the wizard world to Harry and Harry uses some of his new found wealth to buy treats for the train ride.

As Harry meets more of the students, he is a little uncomfortable with his fame. He and Ron meet Hermione Granger, an overachieving girl their age, and Harry sees the unpleasant boy from the uniform shop, Draco Malfoy. When the train stops, Hagrid leads all the first year children to some boats which transport them across a great black lake to the castle of Hogwarts School.

At the door, they are met by Professor McGonagall, who takes them into the Great Hall. Along the way she tells them about the houses they will be separated into, and how the houses earn points with the Quidditch games and can then lose points with infractions. The students are called forward one at a time, a pointed hat is placed on their heads and it announces which house they will be assigned to. The four houses are Gryffindor, Hufflepuff, Ravenclaw, and Slytherin. Harry meets some of the ghosts that reside at the school.

As the sorting ceremony gets under way, Hermione is called forward and the hat puts her in Gryffindor House. Then after a few more, Harry is called. His name is whispered around the room in awe. When the hat is put on his head, it suggests Slytherin,

but, Harry knows that is the house Draco is going into and the house Voldemort was in, so as the hat is put on his head, he keeps thinking, not Slytherin. The hat thinks Slytherin would be a good choice, but complies with Harry's wish and says Gryffindor. When the hat puts Ron in the same house, Harry is thrilled.

After the sorting ceremony is finished, they all settle into eating the sumptuous dinner. While they are eating and getting to know each other, Harry sneaks a glance at the head table where the teachers are all sitting. He notices the teacher of potions, Professor Snape, staring daggers at him. His scar gives a twinge. After dessert, Albus Dumbledore, the Head Master, gives a welcoming speech. He warns them to stay away from the Forbidden Forest and to avoid the third-floor corridor on the right side of the school. They all sing the school song, then retire to their prospective houses.

Harry settles into life at Hogwarts with its moving staircases, paintings that talk and ghosts roaming the halls, conversing with the students. He likes most of his classes but realizes quickly that Professor Snape does not like him. One day Harry receives an invitation for tea at Hagrid's. Harry brings Ron along. At tea, the boys are at first startled by Hagrid's dog, but they discover he is very tame. Hagrid tells Harry not to worry about Snape since he has no reason to dislike Harry. Harry notices a newspaper with an article of a break-in at the bank on Harry's birthday. Harry worries Hagrid may be the culprit.

Flying lessons begin, and Harry is not happy to learn the lessons will be shared with the house of the Slytherins. Whenever Neville hurts his arm and must be led to the infirmary by the instructor, Draco begins to be his usual devilry. He takes a ball that belongs to Neville and takes off with it. Harry mounts his broom in the chase and when Draco tosses it into the air, Harry performs some amazing flying stunts and catches it. Since the instructor told them to stay on the ground, Harry is sure he is in trouble when she pulls him aside, but, instead, she wants him to be on the Quidditch team for the Gryffindor House. Although she wants him to keep it a secret, for now, Harry tells Ron that night at dinner.

Draco and his cronies come up to Harry. After some harsh words between them, Draco challenges Harry to a Wizards Duel. Against Hermione's protests, Harry accepts. That night, when Harry and Ron go to meet Draco, Hermione is still trying to prevent them and gets locked on the wrong side of the door. Now, she and Neville will have to tag along. But, when they get to the meeting place, Draco is nowhere to be seen. Instead, Argus Filch, the school caretaker and his cat, a thoroughly creepy pair, show up. So, the children run for safety. They lose their way and somehow end up on the forbidden third floor, where they are frightened off by a three-headed dog.

After the children have gotten back to their dorm safely, during another one of Hermione's lectures about the proper things to do, Harry notices she says something about the trap door the dog was sitting on. Now, he's curious. Harry gets a new broomstick, and his training as a Quidditch player begins. Meanwhile, during a class on how to make things fly, a skill that Hermione excels at quickly, she overhears a mean comment from Ron and leaves in tears. She doesn't appear at dinner for the Halloween feast and Harry is concerned. Especially when they are told about a troll loose in the school. Realizing that Hermione can't know about the troll, he and Ron search for her.

Before they find Hermione, the two find the troll. They think themselves clever for trapping it in the girls bathroom until they realize that's where Hermione is. Using magic and working together, the three manage to get Hermione free and trap the troll. But, before they can get back to their rooms, Professor McGonagall catches them. As she is reprimanding the boys for leaving their rooms in order to catch the troll, Hermione speaks up that she was the one who was going after the troll and they tried to stop her. After that, she officially becomes their friend.

During Harry's first game of Quidditch, against the Slytherin House, Harry starts having trouble with his broom while trying to catch the Golden Snitch. The team that catches it wins the game. Hermione notices that Professor Snape is staring at Harry and

muttering words, she stops him by catching his robe on fire. Suddenly, Harry has control of his broom again and catches the Golden Snitch, winning the game for Gryffindor.

Later, while having tea at Hagrids shack, Hermione reveals that Snape was controlling Harrys broom. Hagrid disagrees, even when Harry tells him that he thinks Snape was wounded by the three-headed dog, who the children learn is Hagrids dog, Fluffy. He tells them Fluffy is guarding something only Dumbledore and a man named Nicholas Flamel to know about.

During Christmas vacation, Harry has to stay at the school. Although Draco teases him about it, Harry is secretly glad. Christmas with the Dursleys was never very good. The only toys were for Dudley. But, this year, Ron is going to stay there, too, so Harry is quite happy. Christmas morning brings the first gifts he can remember ever receiving. Rons mom knit him a sweater. Also, from a mysterious benefactor he received an invisibility that used to belong to his father. That night after a huge Christmas Dinner, he dons his cloak and goes to the library to explore the restricted books that he has turned away from earlier by the librarian. But, when he opens one of the books, it starts screaming, so he dashes off. Harry ends up in another room where he finds a full-length ornate mirror. When he looks in the mirror he doesnt see himself, but a whole crowd of people. Turning and seeing no one, then when he turns back, he sees his parents, and upon closer inspection, he sees other members of his family. They dont speak to him, they just wave.

The next morning at breakfast, he tells Ron all about the mirror. That night, he and Ron use the cloak for concealment and go back to the mirror. But, instead of his family, Ron sees himself as older, head boy and Quidditch champion. After almost being caught, they rush back to their room. Although Ron warns him not to go back, Harry goes the next night. Dumbledore finds him there. Harry expects him to be angry, but Dumbledore just tells him the mirror shows you your deepest desire and is addictive.

After Christmas break, Hermione returns. The group finally discovers who Nicholas Flamel was. He was Dumbledores partner and the only wizard to have made a Sorcerers Stone. They also learn the stone is supposed to turn lead into gold and make an elixir for immortality. They think this is what the three-headed dog is protecting.

The next Quidditch match Harry catches the Golden Snitch within five minutes and the Gryffindor House takes the championship. Then Harry follows Snape into the woods where he is meeting with Professor Quirrell and talking about the Sorcerers Stone.

As the plot thickens, the children are also worried about homework and exams coming up. Around Easter, Hagrid invites them to see a dragons egg he won in a poker game. He would like to raise it, but they are illegal, so, the children arrange for Rons older brother, who is studying dragons to pick it up after it hatches. Flush with victory, they forget to use the invisibility cloak and are discovered going back to their rooms.

For punishment Ron, Hermione and Harry each lose fifty points for their house. They also have to do detention with Neville and Draco, who were also caught out. They all arrive in the Forbidden Forest where they spend their detention helping Hagrid. Something in the forest has been killing unicorns. After a harrowing time, Harry discovers it was Voldemort drinking unicorn blood to keep going until he can find the Sorcerers Stone.

After learning from Hagrid that he had too much to drink and revealed to Voldemort that the way past Fluffy was with music the children dash off to warn Dumbledore. But, finding him gone, they decide to take the stone and keep it safe. When they get there they see someone else is there already. They get past Fluffy, then fall into a plant that tries to hold them. Hermione uses magic to get them loose. Next is a large room full of birds, that are actually keys. Harrys Quidditch skills come in handy to catch the right bird-key. Then Ron plays a violent game of chess with himself as a piece. He must be captured and beaten by the queen in order to win.

The next hurdle is a series of potions and a logic puzzle. Hermione figures out which potion to drink and then goes back for

Ron. Harry drinks one of the bottles and goes through the flame expecting to find Snape. But, it is Quirrell he finds, instead. Harry discovers that Quirrell has Voldemorts face on the back of his head. Voldemort tells Quirrell to use Harry with the mirror that has been moved into the room. When Harry looks in it, he sees himself putting the stone in his pocket, then it magically is in his pocket. Voldemort tries to get Harry to join him and give him the stone. When he refuses, Voldemort orders Quirrell to capture him. But, he cant touch Harry without burning his hands.

Harry grabs Quirrell and his scar produces excruciating pain. Finally, Harry passes out. When he wakes, Dumbledore is standing over him. Harry tries to tell him that Quirrell took the stone. Harry realizes he is in the medical ward at the school. Dumbledore tells him that he arrived just in time to save him from Quirrell, who couldnt touch him because he was protected by his mothers love. The reason Harry was able to find the Sorcerers Stone is because he was the only one who wanted it for the stone itself, not for what it could do. He also tells him that he and Nicholas Flamel decided to destroy the stone. Then he leaves him with his invisibility cloak.

When Harry meets his friends for the end of year banquet. He is a bit bummed that Slytherin has won the cup, but, then Dumbledore goes up to the podium. He hands out more points. Ron gets fifty points for the best game of chess ever, Hermoine for logic in the face of fire, fifty points. Then for nerve and outstanding courage, Harry was awarded sixty points, and Neville got ten more points for bravery in standing up to his friends. The decorations magically changed to the colors of Gryffindor.

Finally, the grades are handed out, Harry and Ron pass, Hermoine is at the top of her class. After being told not to use magic over the summer break, the children are sailed back across the lake. Then, they board the train back to London. After they go back through the portal, they meet up with their families. Rons little sister, Ginny points out Harry. Harry thanks, Mrs. Weasley for the Christmas gifts and Ron tells him he will get in touch with him over the summer to come over to the Weasleys for a visit. When Hermoine wishes him a good summer, she is looking at the Dursleys uneasily. But, Harry just smiles and says it will be fine, the Dudley doesnt know he isnt allowed to use magic over the summer.

C.3.3 Oliver Twist

Source: <https://www.booksummary.net/oliver-twist-charles-dickens/>

Further away from London a women gave birth to a child and died. Considering the fact that nobody knew a thing about his origin the child got a name and a surname depending on the alphabetical order the children were signed in when they were born. Thats how this child got the name Oliver Twist.

The little kids in this institution were growing up in poverty and misery. The small resources that the institution had were running out because some of the employees were stealing them. Also, the employees used to hit the children and that just increased their misery.

The highest role in the orphanage was played by Mister Bumble, the municipal clerk and the administrator who was in charge

for financing and feeding of the children.

Oliver Twist asked for more porridge once, because he was really hungry, but got beat up at the end. According to the staff, his behavior was rude and that they need to get rid of him as soon as possible. They sent him to a chimney-sweeper so he could learn how to do what he does but he treated him so bad that Oliver was sent back to the orphanage.

After that he came to work for a mortician, Mister Sowerberry. Oliver was a sensitive and pretty boy and he helped the mortician earn a lot by walking next to the childrens coffins. Even though he earned him a lot of money the mortician treated him badly. Oliver slept between the coffins and he would eat only when the dog had some leftovers. Despite the fact that he lived badly with no privileges in life Noa Claypole, the morticians helper and Charlotte, his girlfriend, were jealous of him. The biggest terror over Oliver was done by the landlady. Desperate Oliver attacked the stronger Claypole but he screamed so hard that Oliver was hitting him hard and wants to kill him that Oliver got beat up again.

Oliver decided to put an end to this unbearable situation. He said goodbye to Dick, his dear friend that was ill and dying, and decided to find happiness somewhere else.

Along the way he met a boy named Jack Dawkins. He was using the nickname Artful Dodger. He was a bit weird and rude but Oliver became his friend. Jack took him to the rusty part of London that was filled with thieves, drunkards and wanderers. Oliver met there another weird friend, an old Jewish named Fagin.

Fagin owned a real gang made of grown up criminals, but also little boys. They steal everything that they can and they are no strangers to murders. Artful Dodger alongside a thief named Charles Bates steals in the city. The two of them bring Fagin uppermost of the stolen goods that he sells to other people. Soon he starts to teach Oliver how to steal.

Oliver does not understand that stealing is something bad and sees all of it as a fun game. During a theft the police catches Oliver, even though he did nothing wrong, and brings him to court. He was accused of pickpocketing.

In court he meets Mister Brownlow and he saves Oliver from jail. Oliver lived with him for a while and those were the most beautiful days of his life.

But Oliver's enthusiasm and happiness wears out eventually. Fagin kidnaps the boy with his gang and decides to bring him back to criminal. They did it because a criminal named William Sikes needed a tiny boy to crawl up a window and open the door of the house he wanted to rob. The gang gets help from Nancy, Sikes girlfriend. She sees that Oliver was not made to do this.

Sikes had no mercy. He treated Oliver on the worst possible way and threatened to kill him if he refuses to help him with the robbery. Oliver was forced to do it but the robbery turned out to be a fail at the end. Sikes ran away and left injured Oliver to save himself.

Then Monks comes to the story and he is very mysterious and wants to find Oliver.

Oliver did not know that the robbed house belonged to Maylie and Rose, her adopted daughter and Oliver's aunt. Rose was the sister of his deceased mother. There is another link. Mister Brownlow was a friend of Rose's father.

The two ladies decide to help Oliver and they take care of him after the injury and also prevent him from going to the court again. The boy has special feelings for Rose that even he couldn't explain.

While he was getting better, he met Monks who was constantly following him. In the meantime, Mister Bumble became the manager of the orphanage and Monks visits him in order to find evidence. Monks found out at Mister Bumble that Oliver was his stepbrother. All of that happened thanks to the medallion that was stolen from Oliver's mother when she died.

Sikes hid well after the robbery and his girlfriend Nancy was helping him. Nancy found out why Oliver was so important to everybody. Monks paid Fagin to make Oliver a thief so that he would end up in jail and never receive the fortune his father left

him in his will. The inheritance would then belong to Monks because the two of them shared the same father.

Nancy tells everything to Rose. Fagin finds out what Nancy did and when Sikes hears about it he kills Nancy.

Sikes escapes then. He ran from the police in the company of his dog. Rose was co-operating with Mister Brownlow and he reveals Monks and everything about him. The two of them manage to find a compromise. Monks had to promise that he would not search for Oliver anymore and Brownlow promised not to turn him in.

The police chased Sikes and he died at the end while Fagin and his gang were arrested. The authorities decided to hang Fagin. Mister Bumble stopped being the manager of the orphanage and became a member of it when he ended up homeless.

Mister Brownlow adopted Oliver who then lived a happy and honest life. Oliver's brother Monks moved far away.

C.3.4 The Hound of the Baskervilles

Source: <https://www.booksummary.net/hound-of-baskervilles-arthur-conan-doyle/>

Our story opens at 221B Baker Street, the home and office of one Mr. Sherlock Holmes. His trusty some times assistant, Doctor John Watson, is studying a cane. Dr. Watson has always admired Sherlock's ability to access mountains of information from a small clue, so he thought he would try his hand at it. The cane had been left by the visitor of the night before, it was a fine, thick piece of wood, bulbous-headed, of the sort which is known as a Penang lawyer. Just under the head was a broad silver band nearly an inch across. To James Mortimer, M.R.C.S., from his friends of the C.C.H. was engraved upon it, with the date 1884.

Sherlock, whose back was to Watson, asked what he thought of it. When Watson asked how Sherlock knew what he was doing, he replied that he saw Watson's reflection in the coffee pot. Watson begins to give his deductions, trying to use some of Sherlock's powers of observation. He speculates that the cane belonged to an older doctor, who was given the cane as appreciation for years of service. Getting a bit of encouragement from Sherlock, he continues with the idea that the doctor walked a lot since the cane was well worn. He further reasons that the C.C.H stands for the Something Hunt a local hunt that the doctor must have performed a service.

While Watson is thinking how brilliantly he used his deductive powers, Sherlock tells him it was, interesting, though elementary. Then Sherlock begins to list the observations he has made. While Watson is right about the man being a county practitioner, the C.C.H stands for Charing Cross Hospital. Since the cane would have been given as a going away present, the man must be fairly young, not old as Watson had surmised. And, the man must have a small spaniel dog since there are bite marks on the bottom of the cane. Also, the man and dog are at the door.

The man introduces himself as Mortimer. He is a phrenologist, (he studies skulls to determine intelligence and character). Relieved to see that he had left his cane there, Mortimer proceeds to consult Holmes on a case, since he considers Holmes the second highest expert in Europe. When Holmes asks the first is, the man replies, Monsieur Bertillon. (Doyle was a huge fan of Bertillon who discovered the method of anthropomorphic identification. Before fingerprinting, the exact measurement system Bertillon developed was important in identifying criminals.)

Sherlock suggests that Mortimer go see Bertillon if he's so great. Mortimer replies that Sherlock would be the best choice for this

case, and he hoped he didn't offend. Sherlock says for him to get on with it. Of course, now Sherlock must seem even smarter. So, when Mortimer pulls a manuscript from his pocket, Sherlock tells him, before it's even completely out, that the manuscript is from 1730. The manuscript is actually dated 1742 and tells the story of the curse of the Baskervilles. During the time of the revolution, Hugo Baskerville, who held the manor of Baskerville at the time, was a lecherous, profane and godless man. Hugo desired a local farmer's daughter, who he kidnapped and held her in a room upstairs.

While Hugo was downstairs carousing with his friends, the girl escaped down an ivy-covered wall and set off across the moorland. The furious Hugo made a deal with the devil and released his hounds to chase her. Hugo took off with his dogs and his drunken friends followed afterward. When they found Hugo and the girl, they were both dead. She had died of fear and fatigue, he had had his throat ripped out by some great beast. Since that time the beast has haunted the family of Baskervilles. The most recent victim is Sir Charles Baskerville.

Next, Mortimer shows Sherlock the newspaper article mentioning the death of Sir Charles, a well-respected, philanthropic man, who had remade his family's fortune in South Africa. The story stated that his servants, Mr. and Mrs. Barrymore and Mortimer were all interviewed. According to their reports, Sir Charles was found dead of a heart attack at the site of his nightly walk down Yew Alley, an area bordering the moorlands, where the beast is rumored to roam. The article mentioned the myth of the curse upon the Baskervilles, but only to discount it. The article goes on to name the new owner of Baskerville, Sir Charles' younger brother, Henry, who is in America.

Next, Mortimer tells Sherlock of the facts he left out of the paper. For an unknown reason, Sir Charles had lately become more and more agitated. The curse was playing on his nerves and he thought he saw shadows in the moors. On the night he died, evidence suggested Charles dawdled at the gate to the alley. His footsteps down the alley were curious. He seemed to alternate between tiptoeing and running. But, the most surprising addition, that the newspapers didn't know about, was the footprints of a colossal hound next to Sir Charles' tiptoes.

Now, he has piqued Sherlock's interest. He begins to question Dr. Mortimer. He learns the footprints could not have come from a sheep dog, the paw prints did not approach the body, but were twenty yards away, and the night was damp and raw, but not raining. Then Sherlock wants to know what the alley looks like. He's told there is a strip of grass about six feet wide on either side and it is bordered by an old twelve-foot high yew hedge on each side. The yew is impenetrable. The walk itself is about eight feet across. There is only one opening, a wicket gate leading on to the moor. The exit is through a summer house at the far end, but Sir Charles was within fifty yards of it. The paw prints were on the same side of the path as the moor-gate, and the gate was padlocked, but, it is only four feet high, so anyone could have gotten over it.

After answering all Sherlock's questions, Mortimer said that the reason he hesitated about investigating further, was that, as a man of science, he did not want to entertain the idea of a demon haunting the Baskervilles. But, after interviewing many of the locals, he is left with questions. When Sherlock asks him if he wants him to investigate, Mortimer says that's not why he is there, he just wants advice on what to do about Sir Henry Baskerville, who is due to arrive in an hour from Canada. Sherlock inquires if Sir Henry is the only claimant, and Mortimer tells him that another brother, Rodger, was the black sheep of the family. He fled England to Central America, where he contracted yellow fever and died. There were three brothers; Charles, the second brother, who is Henry's father, and Rodger.

Even though, Mortimer fears for the safety of Sir Henry, he knows the county needs a Lord of the manor at Baskerville to keep the economy moving. Sherlock points out that if the danger is supernatural as Mortimer infers, then he is not safe no matter where he is, so there is no reason to stop him. But, not to tell him of the curse until Sherlock had some time to ruminate on the subject.

He will have an answer the next day, at ten o'clock. Mortimer heads to meet Sir Henry, and Sherlock settles down with his pipe and thinking.

When Dr. Watson returns he finds Sherlock in the room, billowing in smoke. Sherlock surmises Watson went to the club, and then pulls out a map to study the area in question. He wants to discount all the natural possibilities before settling on the supernatural. What would draw a man who was elderly and infirm to go out into the night, wait by the gate for quite a while and, as the change in the footsteps indicate, run, as if for his life, but away from the house, not towards.

Deciding to put aside the thoughts on the mystery until he sees Sir Henry and Mortimer on the next day, Sherlock takes up his violin, to relax. When Sir Henry arrives the next morning, he bears a note that he received warning him away from the Manor House if he valued his life and sanity.

After reviewing the note, Sherlock concludes the note is in a plain envelope with plain rough writing. The note is composed of words cut out of the newspaper, except for the word, moor. Holmes deduces the writer must be following Sir Henry, how else would he know where to find him. The words were cut out of the Times, yesterdays to be exact. Also, the writer used a pair of short nail clippers. Also, the writer must be well educated since only the well educated read the Times. Since the writer was trying to conceal his / her handwriting, their signature must be easily recognizable. And, he must have been in a hurry as the words are glued carelessly. Recognizing he has impressed his audience. Sherlock continues to point out that since the pen was running low on ink, the person is probably staying in a hotel. He tells them that if they try the hotel nearest Charing Cross, they will probably find the rest of the newspaper the words were clipped from.

Sherlock asks Sir Henry if anything unusual has happened and he replied that when he put his boots outside his room for polishing, one was stolen. Holmes decides to tell Sir Henry the story of the curse, but, Sir Henry wants to go to Baskerville anyway. He wants to determine if his uncles death needed a policeman or a clergyman. After inviting Holmes and Watson to lunch later in the day, he and Dr. Mortimer leave.

Sherlock springs to action. He and Watson will grab a cab and follow Sir Henry. They hope to spot the letter writer. When they catch sight of him, they can only make out a bushy black beard. Unfortunately, Sherlock makes a rookie mistake and is spotted by the man. He does manage to note the cab number, though. Sherlock asks a young boy to go through the trash in the hotels in the Charing Cross looking for the cut up Times while he and Watson investigate the cab before they have to meet Sir Henry for lunch. After sending a wire to learn the information on the cab and its occupant, Sherlock and Watson spend some time in a gallery while waiting for their lunch meeting.

When they arrive at the hotel, the desk clerk allows Sherlock to check the registry, where he deduces the spy is not staying in the hotel. When they go upstairs, they find Sir Henry enraged. Another boot has been stolen. This one an older boot. Sir Henry is surprised that Holmes now thinks the thefts may be related to the case.

During lunch, the men discuss the case at length. Holmes discovers that the butler, Mr. Barrymore matches the description of the bearded man. So, Holmes sends a telegram to Baskerville. If Barrymore is not there, it will return to him. Mortimer says that Barrymore inherits 500 pounds and an easy job on Charles death. Also, Mortimer, himself, receives 1000 pounds and Sir Henry 740,000 pounds. The next in line to inherit is a couple of distant cousins, a couple named Desmond. Since Sherlock has prior commitments and cant go to Baskerville right away, he suggests Watson go as a bodyguard to Sir Henry. Before they leave the lunch room, Sir Henry finds his boot.

Back at 221B Bakers Street, Holmes and Watson go over the case but find even fewer answers. Barrymore was not in London, and the cut up newspaper was not located. But, when they interview the cab driver, they think they finally have a clue, until he tells

them the fare used the name, Sherlock Holmes.

When Watson leaves the next day to accompany Sir Henry, Sherlock tells him to only send him the facts and not conjectures, and that he has eliminated the Desmonds as suspects. Then reminds him to bring his gun. Upon arriving at Baskerville, Sir Henry, Mortimer, and Watson see some armed policemen. Apparently, a convict, Seldon, the Notting Hill murderer has escaped and is suspected to be in the area. Mortimer parts ways and heads to his own home. At Baskerville Hall, Sir Henry and Watson are met by Mr. and Mrs. Barrymore. At dinner that night, Sir Henry learns the Barrymores plan to leave his employ and use their inheritance to open a shop. Later, Sir Henry remarks on how creepy the place is, and Watson thinks he hears a woman crying after he goes to bed.

The next morning Watson deduces the crying woman was Mrs. Barrymore and then wonders if perhaps the person in London was Mr. Barrymore after all. So, he questions the post masters boy, only to discover the wire was delivered to Mrs. Barrymore, who said her husband was upstairs. As the mystery deepens, Watson wishes for Sherlock to hurry.

On the way back to Baskerville Watson meets the Stapleton siblings. Brother and sister. He is a naturalist, running around with a butterfly net, and she is beautiful, but a little creepy. Mr. Stapleton walks ways with Watson, giving him a guided tour of the area, saying how glad he is to have Sir Henry there and hoping he is as philanthropic as his uncle. Mr. Stapleton is very nosy, and asks about Sherlock and the case, Watson doesnt comment. While Stapleton runs after a butterfly, Miss Stapleton comes up, to warn Watson to leave. She seems to think he is Sir Henry. But, when she learns otherwise, recants her warning.

In his first report to Sherlock, Watson relates the budding romance between Sir Henry and Miss Stapleton, that her brother doesnt quite approve. He tells of meeting Mr. Frankland, who uses a telescope to search the moors for the escaped convict, who no one has seen for two weeks, so residents think he has left the area. And, also, that he doesnt trust Mr. Barrymore. Sir Henry questioned him about the wire, but he insists he was home. Then there is the candle. One night, Watson follows Barrymore down a hall. He is toting a candle, which he brings to the window, then seems to signal someone outside. Also, Mrs. Barrymore cries every night.

In his second report, Watson speculates that Barrymore is having an affair with a country girl, since the window he signals from faces the moor, and his wife cries so much. When he tells Sir Henry, they decide to stake out the window and confront Barrymore. But, before that can happen, there is a scene between Sir Henry and Miss Stapleton, who he is courting. When he wanted to talk about romance, she kept warning him away from Baskerville, and when he tried to kiss her, her brother came upon them and began to rant. Afterward, he apologized and asked Sir Henry over for dinner next Friday.

After two nights of the stakeout, Watson, and Sir Henry finally, get the chance to catch Barrymore. They find out from Mrs. Barrymore, that the signal is to the escaped convict, who is her brother. They have been bringing him food. Sir Henry and Watson go out to capture the convict for the good of the community. They find him, but he gets away. They hear the moan of a wolf while in the moors, and Watson sees the silhouette of a mysterious tall figure, but it disappears.

Sir Henry and Watson try to convince Barrymore that he needs to help them capture his brother-in-law. But, he asserts the man is harmless and is just waiting to board a ship to South America. They agree to let the matter drop, and as thanks, Barrymore tells them that Sir Charles was supposed to meet a woman the night of his death. They learn the woman was the daughter of Frankland, Laura Lyons. Years ago, she had married a man against her fathers wishes. Her father disowned her and then her husband left her. She is destitute. Sir Charles and Mortimer have been helping her. As for the silhouette, Seldon, the convict has seen him, too. He seems to be a gentleman who lies in a Neolithic hut by the moor. A young boy brings him food. The mystery of the stranger has an entertaining build up, that culminates in finding out the stranger is Sherlock Holmes. He has been staying in the hut, undercover investigation.

He and Watson compare notes. Holmes has discovered that Laura and Mr. Stapleton are close and that Miss Stapleton is actually his wife. Stapleton is actually the villain. He used his wife to lure Sir Henry and Laura. Then he seduced Laura and used her to get to Sir Charles. Sherlock and Watson decide to tell Laura the truth about Stapleton, hoping to change her loyalties. But first, they hear a scream from the moors. At first, they think they have discovered the body of Sir Henry, but, it turns out that Barrymore had given some old clothes of Sir Henrys to Seldon. The hound was given Sir Henrys lost boot to sniff and went for Seldon wearing the clothes. Stapleton arrives thinking to find Sir Henry and is surprised it is not him. Sherlock leads Stapleton to believe he is finished with this case and is going back to London.

After arriving back at Baskerville Hall, Sherlock informs Mrs. Barrymore of the death of her brother and notices a portrait of Hugo Baskerville. He suddenly knows the motive for Stapletons villainy. He is a blood relative of Hugos and obviously just as cruel.

The next morning the plan to prove Stapletons guilt begins. Sherlock tells Sir Henry to keep his appointment for dinner at Stapletons and then walk home through the moors. He also, tells him that he and Watson are going back to London and to tell Stapleton they have left. Then they go to the train station where Sherlock arranges for a wire to be sent from London to Sir Henry, confirming their arrival. Afterward, he and Watson go to Laura and convince her of the truth, that Stapleton is already married and is using her. They learn that Stapleton told her he would marry her and had her contact Sir Charles for help and to meet her. Then had her miss the appointment. Meanwhile, the indomitable Detective Lestrade of Scotland Yard arrives on the scene after being contacted by Sherlock.

Sherlock, Watson and Lestrade spy at the window while Sir Henry and Stapleton are eating dinner. Stapleton goes to a shed and then back inside to his guest. Unfortunately, the weather is not cooperating for following Sir Henry to protect him, but, when he leaves across the moor they manage. When the huge beast comes after him, they all shoot it. Finally, Sherlock empties enough bullets in it to stop it just in time to save Sir Henry. They discover it is a huge dog, mastiff and bloodhound mix. It is as big as a lion and covered with phosphorescent paint. It was truly terrifying enough to frighten Sir Charles to death.

When the detectives go back to Stapletons home, they find Mrs. Stapleton bound and gagged. Upon release she tells them she tried to warn Sir Henry, and where to find her husband. The fog is too thick, still. So they decide to wait for the morning. Meanwhile, they will protect Mrs. Stapleton. The next morning, Mrs. Stapleton leads them on a marked path to her husbands hide out. They find all the proof necessary. But, no Stapleton. Sherlock thinks the swamp killed him.

Back in London, Sir Henry and Mortimer arrive at 221B Baker Street. They ask Sherlock to clarify the mystery to them. He tells them that Stapleton is actually Rodger Baskervilles son. After some legal trouble, he and his wife moved to the area hoping to cash in on his inheritance. He made friends with Sir Charles and discovered his bad heart. Mrs. Stapleton refused to help her husband with his plan, so he romanced Laura. She sent the note and then missed the meeting.

Once Sir Henry arrived on the scene, the Stapletons came to London. Mrs. Stapleton tried to warn Sir Henry, but Stapleton stole his boot for the dog. Unfortunately, the first boot he took was new with no scent, so he had to take another, older boot. Sherlock had suspected the Stapletons from the time of the mysterious note in London because it smelled like perfume.

Mrs. Stapleton did not want to expose her husband until she learned of his relationship with Laura. That was when he tied her up and gagged her. The only loose end Sherlock has is how Stapleton planned on claiming the inheritance. Maybe from South America?

The story ends with Sir Henry accompanying Mortimer on his vacation so he can get some rest. Sherlock and Dr. Watson, plan an outing to dinner and the opera.

C.4 TV Series

C.4.1 Game of Thrones Season 6 Episode 9: Battle of the Bastards

Source: https://en.wikipedia.org/wiki/Battle_of_the_Bastards

In Meereen Daenerys, Tyrion, Missandei and Grey Worm meet with the Masters, who offer to let Daenerys return to Westeros in return for the Masters keeping Missandei and the Unsullied and killing the dragons. Daenerys counters that the meeting was called to discuss the Masters' surrender, and proceeds to ride Drogon into Slaver's Bay with Rhaegal and Viserion to burn their fleet. Missandei tells the Masters that Daenerys has ordered one of them to die as punishment for their crimes. Although they offer the lowborn Yezzan, Grey Worm kills the other two masters instead and Tyrion tells Yezzan to warn the other masters of Daenerys' power. Meanwhile, Daario leads the Dothraki to slaughter the Sons of the Harpy, who are massacring freedmen outside the city.

Theon and Yara arrive in Meereen and offer Daenerys their fleet in exchange for help in overthrowing Euron and recognizing Yara's claim to the Iron Islands. Daenerys agrees to assist them if the Ironborn will stop reaving the mainland, and Yara reluctantly agrees.

At Winterfell Jon, Sansa, Tormund and Davos meet with Ramsay and his bannermen. Ramsay offers to pardon Jon for breaking his vows to the Night's Watch if he hands Sansa over. Jon counters by offering to settle their dispute with single combat; Ramsay refuses, saying that he is more certain that the Bolton army can beat the Stark loyalists than he is of beating Jon one-on-one. When Smalljon Umber proves Rickon's captivity by presenting Shaggydog's head, Sansa tells Ramsay that he will die the next day and rides off. Ramsay gloats that he has been starving his hounds in anticipation of feeding them Jon and his advisors.

After Jon discusses the battle plan with Tormund and Davos, Sansa criticizes him for attacking without gathering more men and predicts that Ramsay will defeat them. Jon insists that their army is the largest one possible. When Jon asks Melisandre not to resurrect him if he dies in battle, she says that it is up to the Lord of Light. Davos and Tormund discuss their time serving Stannis and Mance and acknowledge that they may have served the wrong king all along. Davos discovers the pyre where Shireen and the wooden stag he carved for her were burned.

The armies gather outside Winterfell the next morning. Ramsay brings Rickon out and has him run to Jon while shooting arrows at him. Jon rushes to intercept Rickon, but just before reaching Jon, Rickon is killed by an arrow. Jon charges at Ramsay, who orders the Bolton archers to fire and his cavalry to charge, and Davos orders the Stark force out of position to shield Jon. The ensuing battle leaves hundreds of Bolton and Stark soldiers dead, creating a wall of corpses and allowing the Bolton infantry to encircle the Stark forces. Tormund panics and sends the Wildlings towards the wall of bodies and Smalljon's forces, who easily cut them down. Jon is trampled by the Wildlings, but eventually struggles to his feet. The Stark forces appear doomed when a horn sounds in the distance and Littlefinger and Sansa arrive with the Knights of the Vale, whose cavalry sideswipe and easily smash the remainder of the Bolton army; Tormund kills Smalljon in the chaos.

Ramsay retreats to Winterfell, followed by Jon, Wun Wun and Tormund. Wun Wun breaks down Winterfell's gates and the

Stark loyalists overwhelm the remnants of the Bolton garrison. A mortally wounded Wun Wun is finished off by Ramsay, who tells Jon that he has reconsidered the offer of single combat. Jon blocks Ramsay's arrows with a shield, overpowers him and begins to beat him to death, but stops when he sees Sansa and orders him imprisoned instead, leaving Winterfell once more in the hands of House Stark.

At night, Sansa visits Ramsay, who has been imprisoned in the kennels with his hounds. Ramsay insists that his hounds will not turn on him, but Sansa reminds him that they have been purposefully starved and walks away smiling as they devour him alive.

C.4.2 The Big Bang Theory Season 3 Episode 22: The Staircase Implementation

Source: <https://www.imdb.com/title/tt1648756/plotsummary>

Penny (Kaley Cuoco) is painting her toenails a lovely shade of pink to hide the fact she inherited her father's feet. She has no plans to find out what the argument is between Sheldon (Jim Parsons) and Leonard (Johnny Galecki). Unfortunately, they're so loud she doesn't have to leave her apartment to find out. Leonard asks to sleep on her couch. Leonard asks her if she thinks Sheldon is bat-crazy to argue over the thermostat being set two degrees higher. She says that person is not as crazy as the one who decided to move in with Sheldon. Sounds like a flashback episode, doesn't it? You would be correct in that assumption.

SEVEN YEARS EARLIER

Leonard walks into the building, having borrowed Juan Epstein's hair from Welcome Back, Kotter (1975), and takes the elevator (yes, it works) looking for Sheldon Cooper's apartment. He chose to ignore the warning from Sheldon's former roommate. ("Dude, run away.")

LEONARD: That should have been my first clue.

Penny wondered why that wasn't enough of a warning, but naive Leonard thought his roommate might have been the crazy one instead of Sheldon.

Back in the past, Leonard knocked on the door (Penny's door), but the large, cross-dressing black person he saw wasn't Sheldon. ("Nah, you want the crazy guy across the hall.") Leonard ignored Clue #2 and knocked on Sheldon's door. And it was time for the quiz.

SHELDON: What is the sixth noble gas?

LEONARD: Radon?

SHELDON: Are you asking me or telling me?

LEONARD: Telling you? (pause) Telling you.

SHELDON: Kirk or Picard?

LEONARD: Original Series over Next Generation but Picard over Kirk.

SHELDON: Correct. You've passed the first barrier to roommatehood. You may enter.

Leonard gets to see the apartment, which was bereft of any furniture, save a TV, two lawn chairs, and a bunch of whiteboards. Leonard wants to see the bedrooms, but that would be contingent on making it through the second and third barriers. ("Each more daunting than the last.") Leonard sits down, but not in Sheldon's spot after being rebuked. Sheldon finds out Leonard is an experimental physicist. Sheldon asks if he can drive him, as he chooses not to drive (lie). The final question Leonard asked, but it was easy, since the only thing Sheldon wouldn't want in a post-apocalyptic world is procreation.

LEONARD: Good! I passed the barriers.

SHELDON: Only the second one. Don't get cocky.

It was time for the tour and the first stop was the bathroom.

SHELDON: When do you evacuate your bowels?

LEONARD: Ummm, when I need to?

SHELDON: I'm sorry, but I can't rent to hippies.

LEONARD: OK, 8:00?

SHELDON: I can't give you 8. I can give you 7:30.

Leonard agrees, and he gets to see his room. Granted, Leonard would have to paint over the "DIE SHELDON, DIE!" written in blood on the wall, but it was a minor detail.

Leonard went over the roommate agreement. He is shocked that he is obligated to watch Firefly on Friday nights, but Sheldon thinks it's important since it will be on for years. There is an apartment flag ("don't fly it upside-down unless the apartment is in distress"), and if either invents time travel, the first place they'll go is five seconds from this meeting time. Leonard initials it...and they make sure nobody shows up.

In the present, Leonard defends his decision by saying how nice and affordable the apartment was, and he felt obligated to go through after crashing the first three barriers. And then he brought a girl over.

Leonard is making out with Joyce Kim (big mistake to let her go), and Sheldon initiates his first ****-block. And we learn the origins of the three knock-Leonard cadence. ("I'm just going to keep knocking until you answer.") Apparently, the roommate agreement clearly stated you have to give a coitus warning of 12 hours. (LEONARD: I didn't even KNOW HER 12 hours ago!)

Penny is certainly more sympathetic to Leonard now. (although helping her with her toenails is winning him points in her book). Penny wants to know why he'd stay when Sheldon chased Joyce Kim out of the apartment, but Joyce turned out to be a North Korean spy trying to get some top secret information from Leonard. Keeping Leonard from federal prison was half the reason he stuck with Sheldon. The other half had to do with the elevator.

Back to the past, Sheldon walks in on Leonard...and Raj (Kunal Nayyar) and Howard(Simon Helberg). Seven years ago, Howard was sporting an afro and a soul patch, while Raj apparently didn't get the memo that Miami Vice (1984) was cancelled 10 years earlier. Sheldon was going out of his mind about the leather couch that was now in the apartment. Sheldon thinks he violated the agreement, but Leonard is permitted to decorate 50

PENNY: But what does this have to do with the elevator?

Perhaps it had something to do with Leonard wanting to watch Babylon 5, which Sheldon denies, since nobody in this apartment (meaning him) likes Babylon 5 (1993). But Leonard says it's a tie, and he didn't agree to the "all ties decided by Sheldon" rule. Which Sheldon did agree to, so all ties...

Leonard decides he can fix it, by having everyone leave. Sheldon tries to go with, but he's the guy they're trying to get away

from. ("The correct syntax is: the guy from whom you're trying to get away.") They go to Howard's, and Howard's mother is in good voice tonight. (Didn't sound like Carol Ann Susi, though. Hmmm.) Howard has a three-stage rocket he created, and Leonard thought it was his lucky day because he had rocket fuel back at the apartment. The thing he was going to show Joyce Kim. They return, and Sheldon tells him the rocket fuel won't work because Leonard didn't adjust the formulas correctly to account for a model rocket instead of the real thing. Leonard goes ballistic, and so does the fuel. He tries to take it out of the apartment, but Sheldon does it for him by sealing it in the elevator before it exploded. Sheldon saved his life and didn't rat him out to the police, landlord, or Homeland Security.

PENNY: So the reason I've been walking up and down the stairs for three years is that you did something stupid?

Leonard guffaws, not believing Penny didn't do anything stupid 7 years ago. She was in high school, keeping her nose clean, doing community work, and celebrating the pregnancy test coming back negative.

Back in the present, Leonard apologizes to Sheldon for setting the thermostat to the wrong number. He goes to watch Babylon 5. Oops.

SHELDON: Don't make me turn that flag upside-down, because I'll do it!

C.4.3 Friends Season 5 Episode 14: The One where Everybody Finds out

Source: https://en.wikipedia.org/wiki/The_One_Where_Everybody_Finds_Out

The gang observes that "Ugly Naked Guy," who lives across the street from them, is moving out. Ross (David Schwimmer), who has lived in Joey (Matt Le Blanc) and Chandler's (Matthew Perry) apartment since his botched wedding with Emily, wonders if he should try to get Ugly Naked Guy's apartment. He, Rachel (Jennifer Aniston) and Phoebe (Lisa Kudrow) visit it, and Ross is enthralled, but while he goes for an application, the girls see Chandler and Monica (Courtney Cox) having sex in Monica's apartment. Though initially shocked, Phoebe calms down after Joey and Rachel reveal the two have been together since hooking up at Ross's wedding. Joey, who has been keeping the secret for several months, is relieved that almost everyone knows. However, Rachel and Phoebe want revenge, and decide to mess with the duo by having Phoebe pretend to be attracted to Chandler. Chandler later informs a skeptical Monica that Phoebe was flirting with him.

Upon discovering that Ugly Naked Guy is subletting the apartment himself, Ross attempts to bribe him with a basket of mini-muffins. However, many people have bribed him with extravagant gifts such as a pinball machine and a mountain bike. Ross eventually acquires the apartment after he and Ugly Naked Guy share the mini-muffins whilst nude.

Monica overhears Phoebe flirting with Chandler, and realises he was telling the truth. However, she also realises that Phoebe knows about their relationship and is just trying to mess with them. They confront Joey, who inadvertently reveals Rachel knows as well. Chandler and Monica decide to turn the tables by having Chandler reciprocate Phoebe's advances; to which Rachel and Phoebe realize what the couple are doing and proceed to up the stakes. The game of chicken between the two culminates with Chandler

and Phoebe going on a tense date in Chandler and Joey's apartment. After the two share an awkward kiss, Chandler finally breaks down and reveals he is in love with Monica. Monica (who was hiding in the bathroom) reveals that she is also in love with Chandler, shocking Phoebe who thought they were only in a casual relationship. Joey (who was hiding with Rachel in the hallway) is relieved that he no longer has to keep their relationship a secret. However, the others inform him that they still have to keep it a secret from Ross, much to his chagrin.

In the credits scene, Ross shows his new apartment to his boss, Dr. Ledbetter, to try convince him that he no longer suffers from anger management issues. However, he then sees Monica and Chandler kissing through the window, causing him to angry yell "Get off my sister!"

C.4.4 Narcos Season 1

Source: [https://en.wikipedia.org/wiki/Narcos_\(season_1\)](https://en.wikipedia.org/wiki/Narcos_(season_1))

Season 1 chronicles the life of Pablo Escobar from the late 1970s, when he first began manufacturing cocaine, to July 1992, when he escaped La Catedral prison. The show chronicles the main events that happened in Colombia during this period and Escobar's relationship to them. It is told through the perspective of Steve Murphy, an American DEA agent working in Colombia.

The initial episodes start at the year of 1973 in Chile, with Mateo "Cockroach" Moreno, a Chilean drug dealer and underground chemist in a secret cocaine lab in the Atacama Desert, is discovered by the Chilean Armed Forces and his cartel members are completely executed, but Mateo surprisingly survives, in which later he manages to escape to Colombia. It is followed with the show how Escobar first became involved in the cocaine trade in Colombia. He was an established black marketeer in Medellín, moving trucks worth of illegal goods (alcohol, cigarettes, and household appliances) into Colombia during a time when this was strictly forbidden, when introduced to Mateo, who pitched the idea that they go into business together, with Moreno producing and Escobar distributing a new, profitable drug cocaine. They expand beyond Moreno's small cocaine processing lab by building additional, larger labs in the rainforest and, using the expertise of Carlos Lehder, transport their product in bulk to Miami, where it gains notoriety amongst the rich and famous. Soon enough, Pablo develops larger labs and more extensive distribution routes into the US to supply growing demand. With cocaine's growth into a drug of importance in the American market, one that accounts for a large flow of US dollars to Colombia and escalating drug-related violence in the US, the Americans send a task force from the DEA to Colombia to address the issue. Steve Murphy, the narrator, is partnered with Javier Pea. The role of Murphy's task force is to work with the Colombian authorities to put an end to the flow of cocaine into the United States.

At the time of Murphy's arrival in Colombia, Escobar and his associates are dealing with more significant problems than the DEA. They are at war with the M-19, a revolutionary group of guerilla communists. When the M-19 kidnaps the Ochoa brothers' sister Marta, Escobar seizes the opportunity to form strategic alliances with other black-marketeer criminals to establish a group called "Death to Kidnappers", the genesis of the Medellín cartel. His promise to his allies is simple: to recover Marta Ochoa unharmed and to prevent further kidnappings. In the meantime, Escobar has political aspirations, as he desires to eventually become President of

Colombia. He is elected as a congressman, but is made a fool of when proof of Escobar having criminal ties to the blooming drug industry is brought. Escobar plots his revenge.

An extradition plan is passed in the Colombian congress, allowing for narcos to be extradited to the United States when caught. This is a landmark win for Murphy, Pea, and the DEA and a devastating blow to the Medellín cartel. After making successive threats to the Colombian government to repeal the extradition plan, Escobar takes action against Rodrigo Lara Bonilla, the Colombian Minister of Justice and a prominent lawyer in the prosecution of cartel members, by gunning him down in his car. Murphy and Pea are finally making progress when they catch Escobar's accountant, Blackbeard, along with a gigantic cache of incriminating evidence. The evidence is stored in the only place large enough and thought to have security strong enough to thwart any break-in attempts: the Palace of Justice. However, Escobar hires the M-19, his former enemies, to attack the Palace and burn all of the evidence. The DEA is left with nothing after Escobar's slippery move.

In the sixth episode, Cesar Gaviria the pro-extradition presidential candidate is targeted by Escobar's assassins. Their plan is to blow him up whilst on Avianca Flight 203. He is saved when Murphy successfully warns the candidate not to go on the trip, based on his strong hunch that an assassination attempt is imminent. Nonetheless, with the help of an explosive expert from the terrorist group ETA and with a young, willing, unwitting dupe passenger, the plane is brought down, killing all 107 people on board. The Colombian people are infuriated with the unmanageable levels of violence, especially the plane bombing. The DEA also manages to track down Jos Gonzalo Rodriguez Gacha, one of Escobar's principal associates, and violently gun him and his son Freddy down when they try to escape.

As politicians and his business associates begin to turn against him, Escobar finds a way to strike back at them all: he kidnaps journalist Diana Turbay, the daughter of ex-president Julio Cesar Turbay. Escobar uses Diana as a political bargaining chip to fight the extradition plans that the elected president, Cesar Gaviria, has set in motion, and also to negotiate a peace treaty between the Medellín Cartel and the government. After months of gridlock, the government makes one final attempt to capture Escobar, but government forces mistakenly kill hostage Diana Turbay. As Colombia mourns her death, President Gaviria accepts the terms of Escobar's deal, which will allow Escobar to be incarcerated in his own prison, La Catedral, which will be guarded by his own men. The deal also suspends the extradition plans. Escobar and his colleagues will turn themselves in, and a tentative peace will be restored to Colombia.

Nevertheless, Escobar suffers a tremendous personal blow when his cousin and right-hand man Gustavo Gaviria is brutally killed by the Search Bloc, a designated team of Colombian Special Operations agents tasked with catching high-level drug barons.^[4] In the meantime, Escobar faces competition from the rival Cali Cartel and opposition from members of his own crew. In La Catedral, however, he is protected from authorities and can live in peace without constantly being chased by the Search Bloc. Escobar decorates La Catedral to his liking and hosts guests frequently, including judges, prostitutes, and his family.^[citation needed]

The DEA manages to track the movements of people going in and out of La Catedral and observe that two of Pablo's closest associates, Gerardo Moncada and Fernando Galeano, entered the prison but never left. As the rumors begin to circulate that Escobar killed them in La Catedral, the government uses this new information to try to convince Escobar to be transported to a jail in Bogotá so that the government can further "fortify" La Catedral. Escobar does not approve of this, knowing that once in government hands, he will be prosecuted and extradited to the U.S. The army nevertheless surrounds La Catedral. Eduardo Sandoval, vice-minister of justice, decides to brashly enter the prison to escort Escobar out. He mistakenly assumes he has Pablo's cooperation and is surprised

when Escobar takes him hostage, forcing Sandoval to speak with President Gaviria on Escobar's behalf. President Gaviria, fed up with Escobar's persistent demands, calls for a special-forces team to be sent into the prison to kill Escobar and his men. The team enters, and Sandoval is rescued safely, but Escobar narrowly escapes La Catedral with a few of his men. Outside La Catedral, Escobar is much more vulnerable because he no longer has the protection of his hundreds of men.