

Detecting Narrative Elements in Informational Text

Effi Levi¹, Guy Mor², Tamir Sheaffer^{2,3}, Shaul R. Shenhav²

¹Institute of Computer Science, The Hebrew University of Jerusalem

²Department of Political Science, The Hebrew University of Jerusalem

³Department of Communication and Journalism, The Hebrew University of Jerusalem

efle@cs.huji.ac.il

{guy.mor|tamir.sheafer|shaul.shenhav}@mail.huji.ac.il

Abstract

Automatic extraction of narrative elements from text, combining narrative theories with computational models, has been receiving increasing attention over the last few years. Previous works have utilized the oral narrative theory by Labov and Waletzky to identify various narrative elements in personal stories texts. Instead, we direct our focus to informational texts, specifically news stories.

We introduce NEAT (Narrative Elements Annotation) – a novel NLP task for detecting narrative elements in raw text. For this purpose, we designed a new multi-label narrative annotation scheme, better suited for informational text (e.g. news media), by adapting elements from the narrative theory of Labov and Waletzky (Complication and Resolution) and adding a new narrative element of our own (Success). We then used this scheme to annotate a new dataset of 2,209 sentences, compiled from 46 news articles from various category domains¹. We trained a number of supervised models in several different setups over the annotated dataset to identify the different narrative elements, achieving an average F_1 score of up to 0.77. The results demonstrate the holistic nature of our annotation scheme as well as its robustness to domain category.

1 Introduction

Automatic extraction of narrative elements from texts is a multidisciplinary field of research, combining narrative theories with computational models, which has been receiving increasing attention over the last few years. Examples include modeling narrative structures for story generation (Gervás et al., 2006), using unsupervised methods to detect narrative event chains (Chambers and Jurafsky, 2008) and detecting content zones (Baia-

monte et al., 2016) in news articles, using semantic features to detect *narreme* boundaries in fictitious prose (Delmonte and Marchesini, 2017), identifying turning points in movie plots (Papalampidi et al., 2019) and using temporal word embeddings to analyze the evolution of characters in the context of a narrative plot (Volpetti et al., 2020).

A recent and more specific line of work focuses on using the theory laid out by Labov and Waletzky (1967) and later refined by Labov (2013) to characterize narrative elements in personal experience texts. Swanson et al. (2014) relied on Labov and Waletzky (1967) to annotate a corpus of 50 personal stories from weblogs posts, and tested several models over hand-crafted features to classify clauses into three narrative clause types: *orientation*, *evaluation* and *action*. Ouyang and McKeown (2014) constructed a corpus from 20 oral narratives of personal experience collected by Labov (2013), and utilized logistic regression over hand-crafted features to detect instances of *complicating actions*. More recently, Li et al. (2017) utilized a combination of ideas from Labov and Waletzky (1967) and Freytag (1894) to annotate a collection of short stories, and Saldias and Roy (2020) used convolutional neural networks (CNNs) to classify clauses from spoken personal texts into the same three narrative clause types as Swanson et al. (2014).

While these works concentrated their effort on narrative analysis of personal experience texts, we direct our focus to detecting narrative patterns in informational texts, such as news stories. The social impact of news stories distributed by the media and their role in creating and shaping of public opinion incentivized our efforts to adapt narrative analysis approaches to this domain. To the best of our knowledge, this is the first attempt to automatically detect narrative elements based on Labov and Waletzky (1967) and later works by Labov (1972, 2013) in news articles.

In this work, we introduce NEAT (Narrative Ele-

¹<https://github.com/efle/NEAT>

ments AnnoTation) – a novel NLP task for detecting narrative elements in raw text. For this purpose, we adapted two elements from the narrative theory presented in Labov and Waletzky (1967); Labov (1972, 2013), namely *Complication* and *Resolution*, while adding a new narrative element, *Success*, to create a new multi-label narrative annotation scheme. This scheme was designed with two main objectives in mind. First, capturing elements oriented towards discourse structure, rather than semantic content. Second, possessing the flexibility required to capture narrative characteristics within a wide variety of text types, specifically informational text (as opposed to personal experience), and not only literary and well-structured stories. We used this scheme to annotate a newly-constructed dataset of 2,209 sentences, compiled from 46 English news articles; each sentence was tagged with a subset of the three narrative elements (or, in some cases, none of them), thus defining a novel multi-label classification task.

We explored two different approaches towards solving our new task: splitting into three unrelated binary classification tasks (*Complication*, *Resolution* and *Success*), and jointly learning the three narrative categories as a multi-label classification task. We experimented with three supervised models, each based on fine-tuning a different pre-trained language model: BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019) and DistilBERT (Sanh et al., 2020), achieving an average F_1 score of up to 0.77. An analysis of the results indicates that our narrative categories are strongly connected and form a coherent narrative scheme which is more than just the sum of its parts. Additional experimentation with cross-domain classification demonstrates the task’s robustness to domain category, suggesting that our annotation scheme is more grounded in discourse characteristics rather than semantic context.

The remainder of this paper is organized as follows: Section 2 gives a theoretical background and describes the adjustments we have made to the scheme in Labov (2013) in order to adapt it to informational text. Section 3 provides a complete description of the dataset and of the processes and methodologies which were used to construct and annotate it, along with a short analysis and some examples for annotated sentences. Section 4 describes the experiments conducted on the dataset, and Section 5 provides an analysis and a discus-

sion of the results. Finally, Section 6 contains a summary of our contributions as well as several potential directions for future work.

2 Narrative Analysis

2.1 Background

Ever since the emergence of formalism and structuralistic literary criticism (Propp, 1968) and throughout the development of narratology (Genette, 1980; Fludernik, 2009; Chatman, 1978; Rimmon-Kenan, 2003), narrative structure has been the focus of extensive theoretical and empirical research. While most of these studies were conducted in the context of literary analysis, the interest in narrative structures has made inroads into social sciences (Shenhav, 2015). The classical work by Labov and Waletzky (1967) on oral narratives, as well as later works (Labov, 1972, 2013), signify this stream of research by providing a schema for an overall structure of narratives, according to which a narrative construction encompasses the following building blocks (Labov, 1972, 2013): *abstract* (what is the narrative about), *orientation* (information on the time, the place, the persons and the behavior involved), *complicating action* (or simply complication; the forward progression of narrative clauses), *evaluation* (establishing the narrative’s “point”), *resolution* (what finally happened), and *coda* (bringing the time of reference back to the present time of narration). These building blocks provide useful and influential guidelines for oral narratives analysis.

2.2 Adaptation

Despite the substantial influence of Labov and Waletzky (1967) and Labov (2013), scholars in the field of communication have noticed that this overall structure does not necessarily comply with the form of informational text, such as news stories (Thornborrow and Fitzgerald, 2004; Van Dijk, 1988), and consequently proposed modified narrative structures (Thornborrow and Fitzgerald, 2004). Unlike well-tailored narrative texts, such as personal experience texts, narrativity in informational text is somewhat more challenging as it does not necessarily follow conventional or predefined genre-related structures. This requires a flexible coding scheme, unconstrained by a specific type of text. Instead, it should be open to a wide range of text types (such as informational text), and allow the presence of micro stories, encompassing

	Complication	Resolution	Success
# Sentences	1,092	541	312
Proportion in Dataset	49%	24%	14%

Table 1: Overview of the NEAT dataset. Note that the categories are not mutually exclusive, due to the multi-labeled nature of the annotation scheme.

any combination of all narrative categories even at the sentence level. We set to accomplish that via two objectives: first, formalizing narrative categories which are oriented towards discourse structure, rather than semantic context. Second, defining our task as a multi-labeled one, to allow the flexibility required to capture sentence-level narrative characteristics. A special consideration was given to the variety of contents, forms and writing styles typical for media texts. For example, we required a coding scheme that would fit laconic or problem-driven short reports (too short for full-fledged “Labovian” narrative style), as well as complicated texts with multiple story-lines moving from one story to another. We addressed this challenge by focusing on two of Labov’s six elements - *complicating action* and *resolution*, considered to be the most fundamental and relevant for informational text analysis (Labov, 2013). There are several reasons for our focus on these particular elements: first, it goes in line with the understanding that worth-telling stories usually consist of protagonists facing and resolving problematic experiences (Eggin and Slade, 2005). Moreover, these elements resonate with what is considered by Entman (2004) to be the most important Framing Functions - problem definition and remedy.

In order to adapt the original *complicating action* and *resolution* categories to informational content, we designed our annotation scheme as follows. *Complicating action* – hence, *Complication* – was defined in our narrative scheme as an event, series of events or situation, that point at problems or tensions. *Resolution* refers to the way the story is resolved or to the release of the tension. An improvement from – or a manner of coping with – an existing or a hypothetical situation was also considered to be a *Resolution*. This choice was made in order to follow the often tentative or speculative notion of future resolutions in news stories (Thornborrow and Fitzgerald, 2004; Bell, 1991). We have therefore included in this category any temporary or partial resolutions. The transitional characteristic of the *Resolution* motivated us

to add a new category defined as *Success*. Unlike *Resolution*, which refers, implicitly or explicitly, to a prior situation, this category was designed to capture any description or indication of an achievement or a desirable outcome.

3 The Dataset

3.1 Pilot Study

We started by conducting a pilot study, for the purpose of formalizing an annotation scheme and training our annotators. For this study, sample sentences were gathered from print news articles, published between 1995 and 2017 and collected via *Lexis-Nexis*. These were used to refine the annotation scheme described in Section 2.2, as well as perform extensive training for our annotators.

Following the conclusion of the pilot study, we used the sentences which were collected and manually annotated during the pilot to train a multi-label classifier, later used to provide labeled candidates for the annotators during the annotation stage of the NEAT dataset, in order to optimize annotation rate and accuracy. The pilot samples were then discarded.

3.2 News Articles

The news articles for the dataset were sampled from leading news websites in the English language, all published between 2017 and 2020. The result is a corpus of 2,209 sentences taken from 46 news articles, with an average of 48 sentences per article ($\sigma^2 = 39.44$), and an average of 20.2 tokens per sentence ($\sigma^2 = 11.2$). The articles are semantically diverse, as they were sampled from a wide array of domain categories.

3.3 Preprocessing

The news articles’ content was extracted using *diffbot*. The texts were scraped and split into sentences using the Punkt unsupervised sentence segmenter (Kiss and Strunk, 2006). Remaining segmentation errors were manually corrected.

3.4 Annotation

3.4.1 Guidelines

Following the pilot study (Section 3.1), a code book containing annotation guidelines was produced. For each of the three categories in the annotation scheme – *Complication*, *Resolution* and *Success* – the guidelines provide:

- A general explanation of the category
- A list of well-defined criteria for identifying the category
- Select examples of sentences labeled with the category

3.4.2 Process

We employed a three-annotator setup for annotating the collected sentences. First, the pilot stage model (Section 3.1) was used to produce annotation suggestions for each of the sentences in the corpus. Each sentence was then separately annotated by two trained annotators according to the guidelines described in Section 3.4.1. Each annotator had the choice to either accept the suggested annotation or to change it by adding or removing any of the suggested labels. Disagreements were later decided by a third expert annotator. Table 2 reports inter-coder reliability scores for each of the three categories, averaged across pairs of annotators: pairwise percent agreement (PPA), and Cohen’s Kappa coefficient, accounting for chance agreement (Artstein and Poesio, 2008). Article-level domain categories (Table 3) were initially assigned according to the news section from which the articles were taken, and later verified by two annotators.

	PPA (avg.)	κ (avg.)
Complication	89.6%	0.79
Resolution	87.5%	0.65
Success	92.8%	0.7

Table 2: Inter-coder reliability: average pairwise percent agreement (PPA) & average Cohen’s Kappa (κ)

3.5 Analysis

Narrative categories vary significantly in their prevalence in the corpus; their respective proportions in the dataset are given in Table 1. The categories are unevenly distributed: *Complication* is significantly more frequent than *Resolution* and *Success*. This was to be expected, considering the known biases of "newsworthiness" towards

problems, crises and scandals (Esser et al., 2016), and due to the fact that in news media, resolutions often follow reported complications.

Interestingly, the distribution over narrative categories varies significantly between the different category domains (see Table 3). Most domains contain many more *Complications* than *Resolutions* or *Successes*, which is consistent with the distribution in the complete dataset (Table 1); the “Crime” domain is an extreme example, with a very small number of *Resolutions* and no *Successes* at all. However, some domains exhibit a completely different distribution. For example, the “Travel” and the “Science & Technology” domains possess a relatively uniform distribution over the three narrative categories. The “Sports” domain contains a similar number of *Complications* and *Successes*, with a smaller number of *Resolutions*.

Table 4 reports pairwise Pearson correlations between the categories. The *Complication* and *Resolution* categories are completely uncorrelated ($r = 0.016$). A minor negative correlation was found between *Complication* and *Success* ($r = -0.234$), and a minor positive one was found between *Resolution* and *Success* ($r = 0.228$). These minor correlations – in our opinion – indicate that the *Success* category does indeed bring added value to our narrative scheme.

All the possible combinations of narrative categories appear in the dataset; Table 5 summarizes the occurrences of each of the possible category combinations. Examples of sentences annotated with various category combinations are given in Appendix A. There is, however, a significant variability to the frequency in which different combinations occur in the dataset. For example, the *Complication-Resolution* combination, designating a typical narrative tension-relief pattern (Shenhav, 2015), is by far the most frequent one with 226 sentences. *Complication-Success*, on the other hand, is a very rare combination with only 15 sentences, embodying a far less trivial or common logic, where a success is accompanied by an unresolved problem.

The fact that the dataset is assembled from full coherent news articles allows the analysis of a range of micro and macro stories in narrative texts. For example, an article in the dataset concerning

Domain Category	# Sentences	Complication	Resolution	Success
Economy & Job Market	221	99	71	25
Politics	860	514	220	89
Health	66	38	13	10
Travel	200	47	46	42
Arts & Culture	345	166	63	49
Crime	28	25	3	0
Accidents & Disasters	12	9	4	2
Welfare	47	30	11	1
Sports	115	51	21	47
Science & Technology	223	53	60	40
Immigration	92	60	29	7

Table 3: Domain category distribution in the NEAT dataset

	Comp.	Res.	Suc.
Comp.	1		
Res.	0.016	1	
Suc.	-0.234	0.228	1

Table 4: Inter-category Pearson correlations

Combination	# Sentences
Complication & Resolution	226
Complication & Success	15
Resolution & Success	103
All Three Categories	49

Table 5: Narrative category co-occurrences in the NEAT dataset

the coronavirus outbreak in South Korea² opens with a one-sentence summary, tagged with both *Complication* and *Resolution*:

"South Korea's top public health official hopes that the country has already gone through the worst of the novel coronavirus outbreak that has infected thousands inside the country." (*Complication, Resolution*)

This problem-solution (in this case, hopeful solution) plot structure reappears in the article, this time detailed over a series of sentences:

"More than 7,300 coronavirus infections have been confirmed throughout South Korea, killing more than 50." (*Complication*)

"It is one of the largest outbreaks outside mainland China, where the deadly virus was first identified." (*Complication*)

²<https://edition.cnn.com/2020/03/09/asia/south-korea-coronavirus-intl-hnk/index.html>

"However, the number of new daily infections in South Korea has declined in recent days." (*Complication, Resolution*)

"... while he believes the aggregate number of infections is high, he is confident in the job South Korea did to combat the virus' spread and would advise other governments..." (*Complication, Resolution*)

"The South Korean government has been among the most ambitious when it comes to providing the public with free and easy testing options." (*Success*)

The sequence starts with two sentences tagged with *Complication*, followed by two sentences tagged with *Complication* and *Resolution*, and concludes with a sentence tagged with *Success*, demonstrating a more gradual transition from problem through solution to success.

4 Experiments

4.1 Dataset Partition

We randomly divided the dataset into article-wise mutually-exclusive train, validation and test sets (details in Table 6), while keeping the distribution over the three narrative categories in each of the sets as similar as possible to the one in the complete dataset. The train set was used to train a supervised model for the task; the validation set was used to select the best model configuration during the training phase by tuning the model's hyper-parameters (see Section 4.3 for details), and the test set was used to evaluate the chosen model and produce the results reported in Section 5.

	# Sentences (Articles)	Ratio
Train	1,767 (37)	80%
Validation	222 (5)	10%
Test	220 (4)	10%

Table 6: Train, validation & test set statistics

4.2 Task Definition

We explored two different approaches for solving the task: (1) addressing each of the three narrative categories as a separate classification task, and (2) treating the task as multi-label classification with three labels (one for each narrative category).

4.2.1 Separate Classification Tasks

In this approach, we defined a separate binary classification task for each of the narrative categories: *Complication*, *Resolution* and *Success*. For each such task, we trained a dedicated supervised model (further details given in Section 4.3), specifically optimized for the learned category. However, any potential information stemming from inter-correlations between the different categories was ignored, effectively treating them as three unrelated tagging schemes.

4.2.2 Multi-Label Classification

Here, the task was treated as a three-way multi-label classification problem (each sentence may contain any combination of the three narrative categories), thus taking advantage of inter-correlations between the three narrative categories to better learn them as part of a coherent narrative scheme. We trained and optimized a single multi-label model to jointly predict the three categories (further details given in Section 4.3).

4.3 System Architecture

We employed the method of fine-tuning a pre-trained language model for our task. In each experimental setup, we chose a pre-trained language model as a backbone, applied a multilayer perceptron (MLP) classifier on top of it, and fine-tuned the entire model over the train set.

4.3.1 Backbone

We experimented with three state-of-the-art transformer-based language models as the backbone for our inference model, using pre-trained weights from the *transformers* python package (Wolf et al., 2019).

BERT. Following common practice, we first utilized the base-sized BERT (Devlin et al., 2018) as the backbone model.

RoBERTa. This BERT variant was developed by training the original BERT model with altered design choices and training techniques, and has been recently shown to produce better results on various NLP tasks (Liu et al., 2019). We used the base-sized RoBERTa as the backbone model.

DistilBERT. A recent body of work has focused on developing “lighter” transformer-based language models which allow for faster fine-tuning for downstream NLP tasks (Sanh et al., 2020; Lan et al., 2020). In order to assess robustness to a decrease in the model’s size, we also experimented with DistilBERT (Sanh et al., 2020), which follows the same basic architecture as BERT but consists of 66M parameters (as opposed to 110M in BERT), as the backbone model.

4.3.2 Classifier

In order to fine-tune the backbone language model, we appended a multilayer perceptron (MLP) over the output of the language model. The MLP consisted of one hidden layer (increasing the number of hidden layers produced no improvement in performance), the size of which was optimized as a hyper-parameter. In the case of a single binary classification task (Section 4.2.1), the output layer consisted of a single sigmoid output, while in the case of a multi-labeled task, it consisted of three sigmoid outputs, one for each narrative category.

4.3.3 Training Procedure

All models were optimized using the AdamW algorithm (Loshchilov and Hutter, 2017) and the binary cross entropy loss function. Positive weighting was used in order to compensate for class imbalance (evident in Table 1). Hyper-parameters – batch size, learning rate and MLP hidden layer size – were chosen via a standard grid search (see Appendix B for more details). For each configuration of task definition, backbone model and hyper-parameters, the model was evaluated over the validation set after every epoch of training, and the best-performing checkpoint was tested on the test set to produce the results reported in Section 5.

4.4 Cross-Domain Classification

Given the semantic diversity in the dataset, as well as the variability in distribution over the narrative categories between the various domains (Table 3),

Task	Backbone	Complication			Resolution			Success			Average		
		P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
Separate	BERT	0.89	0.87	0.88	0.67	0.40	0.50	0.33	0.72	0.45	0.63	0.66	0.61
Separate	RoBERTa	0.92	0.85	0.88	0.69	0.56	0.62	0.83	0.56	0.67	0.81	0.66	0.72
Separate	DistilBERT	0.87	0.84	0.85	0.69	0.40	0.51	0.53	0.44	0.48	0.70	0.56	0.61
M.Label	BERT	0.88	0.89	0.88	0.72	0.60	0.65	0.53	0.50	0.51	0.71	0.66	0.68
M.Label	RoBERTa	0.90	0.90	0.90	0.76	0.67	0.71	0.64	0.78	0.70	0.77	0.78	0.77
M.Label	DistilBERT	0.86	0.87	0.86	0.73	0.53	0.61	0.75	0.50	0.60	0.77	0.63	0.69

Table 7: Test set precision (P), recall (R) and F_1 scores, for every combination of task definition (a separate task for each narrative category / a multi-labeled task) and backbone model. See Sections 4.2 and 4.3.1 for details.

we wished to assess the domain category’s effect on learning our narrative scheme. For this purpose, we experimented with a cross-domain classification setup. For each of the eleven category domains, we concatenated the sentences from all other domains into a train set, which was then used to train a classification model. The training process was done using the configuration of the best-performing model from the previous stage (described in Sections 4.1–4.3), including task definition, backbone model and hyper-parameters (i.e. no hyper-parameter tuning was performed in this setup). The trained model was then evaluated on the test set.

5 Results & Discussion

Results are reported in Table 7. For each task definition and choice of backbone model, we report the precision, recall and F_1 score for each of the three narrative categories, as well as their average, over the test set.

First, we observe that addressing the task as a multi-labeled one proved to be a better strategy than learning each narrative category separately. This is evident across backbone models as well as across narrative categories; for each backbone model, the multi-label model produced a higher F_1 score for each and every one of the narrative categories. This is a clear indication that these categories are substantially connected as they constitute intertwining elements in an underlying story. Therefore, the three categories form a coherent narrative scheme that is more than just the sum of its parts.

Interestingly, while this effect is relatively small for *Complication* (F_1 increased by 0.00–0.02), it is much more prominent for *Resolution* (F_1 increased by 0.09–0.15) and *Success* (F_1 increased by 0.03–0.12), meaning that incorporating all three categories into one coherent scheme contributes mostly to learning the *Resolution* and *Success* categories. This

suggests that perhaps the narrative properties of the *Complication* category make it more independent and self-contained than the other two categories. *Resolution* and *Success*, on the other hand, are more relative in nature, and seem to be anchored, implicitly or explicitly, by a prior situation or condition.

Among the three narrative categories, *Complication* gained the highest F_1 scores by all the models, ranging between 0.85 and 0.90. The models were less successful predicting *Resolution*, with F_1 scores ranging between 0.50 and 0.71, and *Success*, with F_1 scores ranging between 0.45 and 0.70. This is consistent with the proportion of instances belonging to each category in the dataset (see Table 1), which may provide a possible explanation for this observation. However, the fact that positive weighting was used in the training process (Section 4.3.3) to counter class imbalance, motivates us to search for another explanation. Defined by Labov and Waletzky (1967)’s overall structure of narratives as “the main body of narrative clause”, *Complication* may just be an easier narrative category to learn.

Comparing different backbone models, the DistilBERT-based model performed similarly to the BERT-based one – an average F_1 score of 0.61 for both in the separate-task setting, and 0.69 compared to 0.68 in the multi-label setting – suggesting that the task is fairly robust to a decrease in the backbone model’s size. However, RoBERTa significantly outperformed the other two language models as the backbone in both settings – an average F_1 score of 0.71 compared to 0.61 in the separate-task setting, and an average F_1 score of 0.77 compared to 0.68 and 0.69 in the multi-label setting. We also note that the difference in performance between *Complication* and the other two categories is less extreme in the RoBERTa-based models compared to the other backbone models.

Domain Category	Complication			Resolution			Success			Average		
	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
Accidents & Disasters	1.00	0.89	0.94	1.00	0.50	0.67	1.00	0.50	0.67	1.00	0.63	0.76
Arts & Culture	0.83	0.84	0.84	0.67	0.75	0.71	0.51	0.78	0.62	0.67	0.79	0.72
Crime	1.00	0.92	0.96	0.67	0.67	0.67	-	-	-	0.83	0.79	0.81
Economy & Job Market	0.86	0.78	0.81	0.77	0.68	0.72	0.50	0.84	0.63	0.71	0.76	0.72
Health	0.97	0.82	0.89	0.63	0.77	0.69	0.67	0.80	0.73	0.75	0.80	0.77
Immigration	0.86	0.93	0.90	0.80	0.83	0.81	0.71	0.71	0.71	0.79	0.83	0.81
Politics	0.83	0.95	0.89	0.58	0.86	0.69	0.60	0.65	0.63	0.67	0.82	0.74
Science & Technology	0.75	0.83	0.79	0.66	0.77	0.71	0.64	0.88	0.74	0.68	0.82	0.74
Sports	0.87	0.65	0.74	0.53	0.76	0.63	0.84	0.81	0.83	0.75	0.74	0.73
Travel	0.89	0.87	0.88	0.66	0.89	0.76	0.69	0.83	0.75	0.75	0.87	0.80
Welfare	1.00	0.97	0.98	0.80	0.73	0.76	0.33	1.00	0.50	0.71	0.90	0.75
Average	0.90	0.86	0.87	0.71	0.74	0.71	0.65	0.78	0.68	0.76	0.79	0.76

Table 8: Precision (P), recall (R) and F_1 scores for cross-domain classification. See Section 4.4 for details.

5.1 Cross-Domain Classification

The best-performing configuration (a multi-label classifier based on the RoBERTa language model) was used to perform the cross-domain classification experiment. As stated in Section 4.4, hyperparameters were fixed to the values obtained in the train-validation-test setup. Results are presented in Table 8. Averaged over domain categories, they are virtually identical to the results obtained on the test set (reported in Table 7), with a precision, recall and F_1 score of 0.76, 0.79 and 0.76 compared to 0.77, 0.78 and 0.77 (respectively). In our opinion, this demonstrated invariance to domain category is a strong indication that our narrative elements are more grounded in discourse characteristics rather than in the semantic field.

6 Conclusion

We introduced NEAT (Narrative Elements Annotation) - a novel NLP task for detecting narrative elements in raw text. For this purpose, we designed a new flexible multi-label narrative annotation scheme, specifically suited for informational text, by adapting two elements from the theory introduced in Labov and Waletzky (1967); Labov (1972, 2013) - *Complication* and *Resolution* - and adding a new element - *Success*. The scheme was used to annotate a new dataset of 2,209 sentences, compiled from 46 articles, which were collected from news websites.

We explored two alternate settings for solving this task - one in which each narrative category was treated as a separate classification task, and another in which the entire task was addressed as multi-label classification. In each of these setups,

we experimented with fine-tuning three different language models, achieving an average F_1 score of up to 0.77 on the test set, and showcasing the potential of supervised-learning methods in detecting the narrative information encoded into our scheme. The multi-label setting consistently provided significantly better results across all models and narrative categories, demonstrating that our narrative categories are strongly connected and form a coherent narrative scheme which is more than just the sum of its parts. Additional cross-domain classification results demonstrate the task’s invariance to domain category, suggesting that our annotation scheme is more grounded in discourse characteristics rather than semantic context.

We are currently engaged in an ongoing effort for improving the annotation quality of the dataset and increasing its size. In addition, we have several interesting directions for future work. The first one, which we are currently pursuing, includes enriching the scheme with token-level annotation of the narrative elements, effectively converting the task from multi-label classification to a sequence prediction one. Alternatively, we could introduce additional layers of information to encode more global narrative structures in the text, such as inter-sentence - or even inter-article - references between narratively-related elements (e.g., a *Resolution* referencing its inducing *Complication*). Another potential direction is incorporating additional narrative elements into our annotation scheme. For example, the *evaluation* element from (Labov, 2013) may be beneficial in encoding additional information in the context of news media, such as the sever-

ity of a Complication or the ‘finality’ of a Resolution. We could also add completely new narrative elements, tailored to capture specific informational aspects, such as actor-based elements identifying entities which are related to one or more of the currently defined narrative categories.

Acknowledgements

This research was supported by the Israel Science Foundation (Grants No. 1400/14; 2315/18). We thank the three annotators – Gal Ron, Vered Porzycki and Avishai Green – for their dedicated work, as well as their insightful comments and suggestions during the annotation processes.

References

- Ron Artstein and Massimo Poesio. 2008. [Survey article: Inter-coder agreement for computational linguistics](#). *Computational Linguistics*, 34(4):555–596.
- Daniela Baiamonte, Tommaso Caselli, and Irina Prodanof. 2016. Annotating content zones in news articles. *CLiC it*, page 40.
- Allan Bell. 1991. *The language of news media*. Blackwell Oxford.
- Nathanael Chambers and Dan Jurafsky. 2008. [Unsupervised learning of narrative event chains](#). In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.
- Seymour B Chatman. 1978. Existents. *Story and Discourse: Narrative Structure in Fiction and Film*, pages 96–145.
- Rodolfo Delmonte and Giulia Marchesini. 2017. [A semantically-based computational approach to narrative structure](#). In *IWCS 2017 — 12th International Conference on Computational Semantics — Short papers*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Suzanne Eggins and Diana Slade. 2005. *Analysing casual conversation*. Equinox Publishing Ltd.
- Robert M Entman. 2004. *Projections of power: Framing news, public opinion, and US foreign policy*. University of Chicago Press.
- Frank Esser, Sven Engesser, and Jörg Matthes. 2016. Negativity. In *Comparing political journalism*, pages 89–109. Routledge.
- Monika Fludernik. 2009. *An introduction to narratology*. Routledge.
- Gustav Freytag. 1894. *Die technik des dramas*. S. Hirzel.
- G rard Genette. 1980. Narrative discourse, tr. J. Lewin. *Ithaca: Cornell [= 1972, Figures III*.
- Pablo Gerv s, Birte L nneker-Rodman, Jan Christoph Meister, and Federico Peinado. 2006. Narrative models: Narratology meets artificial intelligence. In *International Conference on Language Resources and Evaluation. Satellite Workshop: Toward Computational Models of Literary Analysis*, pages 44–51.
- Tibor Kiss and Jan Strunk. 2006. [Unsupervised multi-lingual sentence boundary detection](#). *Computational Linguistics*, 32(4):485–525.
- William Labov. 1972. *Language in the inner city: Studies in the Black English vernacular*, volume 3. University of Pennsylvania Press.
- William Labov. 2013. *The language of life and death: The transformation of experience in oral narrative*. Cambridge University Press.
- William Labov and Joshua Waletzky. 1967. Narrative analysis. oral versions of personal experience. *Essays on the Verbal and Visual Arts: Proceedings of the American Ethnological Society*, edited by June Helm, 12-44. Seattle: American Ethnological Society.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#).
- Boyang Li, Beth Cardier, Tong Wang, and Florian Metze. 2017. Annotating high-level structures of short stories and personal anecdotes. *arXiv preprint arXiv:1710.06917*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Jessica Ouyang and Kathy McKeown. 2014. [Towards automatic detection of narrative structure](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*, pages 4624–4631, Reykjavik, Iceland. European Languages Resources Association (ELRA).
- Pinelopi Papalampidi, Frank Keller, and Mirella Lapata. 2019. Movie plot analysis via turning point identification. *arXiv preprint arXiv:1908.10328*.
- Vladimir Yakovlevich Propp. 1968. *Morphology of the Folktale*, volume 9. University of Texas Press.
- Shlomith Rimmon-Kenan. 2003. *Narrative fiction: Contemporary poetics*. Routledge.

- Belen Saldias and Deb Roy. 2020. [Exploring aspects of similarity between spoken personal narratives by disentangling them into narrative clause types.](#)
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.](#)
- Shaul R Shenhav. 2015. *Analyzing social narratives.* Routledge.
- Reid Swanson, Elahe Rahimtoroghi, Thomas Corcoran, and Marilyn Walker. 2014. [Identifying narrative clause types in personal stories.](#) In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 171–180, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Joanna Thornborrow and Richard Fitzgerald. 2004. Storying the news through category, action, and reason. *The Communication Review*, 7(4):345–352.
- Teun A Van Dijk. 1988. *News as discourse.* hillsdale, nj: L.
- Claudia Volpetti, K Vani, and Alessandro Antonucci. 2020. Temporal word embeddings for narrative understanding. In *Proceedings of the 2020 12th International Conference on Machine Learning and Computing*, pages 68–72.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing.](#)

Appendix

A Sample Annotated Sentences

1. How did some of the biggest brands in care delivery lose this much money? (Complication)
2. Bleeding from the eyes and ears is also possible after use, IDPH said. (Complication)
3. His proposal to separate himself from his business would have him continue to own his company, with his sons in charge. (Resolution)
4. Instead, hospitals are pursuing strategies of market concentration. (Resolution)
5. With its centuries-old canals, vibrant historic center and flourishing art scene, Amsterdam takes pride in its cultural riches. (Success)
6. Mr. Trump chose to run for president, he won and is about to assume office as the most powerful man in the world. (Success)
7. Soon after, her administration announced a set of measures intended to curb misconduct. (Complication, Resolution)
8. Avant-gardists, the couple opened an art gallery in 1875 within the department store, offering artists who had been turned away by the Paris Salon – the official art exhibition of the Academy of Fine Arts in Paris – a home for their works and a large public audience, Burckhardt wrote. (Complication, Resolution)
9. Unlike Macy’s, the well-known US department store which has been closing its doors around the country, Le Bon Marché isn’t fighting to stay alive. (Complication, Success)
10. He defeated Bolt in a close semifinal heat, ending Bolt’s 28-race winning streak in the 100 meters. (Complication, Success)
11. The Utah man’s mother, Laurie Holt, thanked Mr. Trump and the lawmakers for her son’s safe return, adding: "I also want to say thank you to President Maduro for releasing Josh and letting him to come home." (Resolution, Success)

12. And if Tony Hughes can keep his job without the weeks away, "in that aspect," he said, "it’s going to make my life better." (Resolution, Success)
13. They were fortunate to escape to America and to make good lives here, but we lost family in Kristallnacht. (Complication, Resolution, Success)
14. He was vulnerable, no doubt, but in the past, with the spotlight at its brightest, he still found the speed and the will to remain the fastest man in the world. (Complication, Resolution, Success)

B Hyper-Parameters

Table 9 lists the ranges of values which were used for hyper-parameter tuning during the experiments described in Section 4.

Hyper-Parameter	Values
Batch Size	16, 32, 64
Learning Rate	1e-5, 2e-5, 3e-5, 4e-5, 5e-5
MLP Hidden Size	50, 100, 200, 300

Table 9: Hyper-parameters value ranges

The best-performing model – a multi-label classifier based on the RoBERTa language model – was trained for 4 epochs with a batch size of 32, a learning rate of 4e-5 and an MLP classifier with a 100-node hidden layer.