

# Saliency-Aware Event Chain Modeling for Narrative Understanding

Xiyang Zhang, Muhao Chen, and Jonathan May

Information Sciences Institute

University of Southern California

{xiyangzh, muhao, jonmay}@isi.edu

## Abstract

Storytelling, whether via fables, news reports, documentaries, or memoirs, can be thought of as the communication of interesting and related events that, taken together, form a concrete process. It is desirable to extract the event chains that represent such processes. However, this extraction remains a challenging problem. We posit that this is due to the nature of the texts from which chains are discovered. Natural language text interleaves a narrative of concrete, salient events with background information, contextualization, opinion, and other elements that are important for a variety of necessary discourse and pragmatics acts but are not part of the principal chain of events being communicated. We introduce methods for extracting this principal chain from natural language text, by filtering away non-salient events and supportive sentences. We demonstrate the effectiveness of our methods at isolating critical event chains by comparing their effect on downstream tasks. We show that by pre-training large language models on our extracted chains, we obtain improvements in two tasks that benefit from a clear understanding of event chains: narrative prediction and event-based temporal question answering. The demonstrated improvements and ablative studies confirm that our extraction method isolates critical event chains.<sup>1</sup>

## 1 Introduction

Human languages always communicate about evolving events. Hence, it is important for NLP systems to understand how events are procedurally described in text. In this context, identifying patterns of event chains is important but challenging, as it requires knowledge of inter-event relations such as temporal or causal relations. Modeling high-quality event chain patterns from text would be a first step toward the more general goal of

<sup>1</sup>Our code and data are available at <https://github.com/juvezxy/Saliency-Event-Chain>.




Figure 1: An illustration of event chains and saliency-aware and discourse-aware filtering on an example text. The words in *italic* are the events and the arrows between them show the temporal relation given by TEAR (pointing from previous to next). Salient events are shown in **bold**; the blue sentence is in the *Main Event* category whereas the red sentence is in the *Evaluation* category.

event schema induction, which involves generating high-level representations of event relations and structures. The extracted chain patterns can in turn represent useful information for core natural language tasks such as question answering (Reddy et al., 2019), semantic role labeling (Cheng and Erk, 2018), story generation (Yao et al., 2019), and reading comprehension (Ostermann et al., 2019).

Generally speaking, “events” correspond to what we perceive as happening around us. According to the theory of embodied cognition (Wilson, 2002), our understanding of the world is shaped by various aspects of our entire bodies, involving also our language comprehension. However, it is currently difficult to gather enough data from other modalities to model real world “events,” and written text, especially in the news domain, seems to be our best option.

Previous attempts have been made to generate event chains by modeling narratives in news, stories and documentaries (Chambers and Jurafsky, 2008; Weber et al., 2018b; Li et al., 2020). The problem with such data is that the articles are usually a mixture of the true narrative flow with other content, which serves to explain the context or provide side information. Most prior approaches do not take

into account the centrality or salience of events or the discourse structures that describe those events. Accordingly, overlooking such important discourse properties when choosing the events of a sequence introduces noise in the modeling of event chains, and causes trivial or irrelevant content to be captured and inferred in narrative understanding tasks.

In this work, we explore the use of salience identification (Liu et al., 2018; Jindal et al., 2020) and discourse profiling (Choubey et al., 2020) to help isolate the main event chains from other distracting events, and show the effect on two recent datasets related to narrative understanding and temporal understanding. More specifically, we obtain event chains from documents and perform different methods of filtering, then build event language models, which we use to predict story ending events from the ROCStories dataset (Mostafazadeh et al., 2016) and answer temporal event questions from the TORQUE dataset (Ning et al., 2020). By comparing the use of event language models built on differently filtered event sources, we show that filtering out distracting narrative information can indeed benefit the modeling of event relations, thus leading to more reliable extraction of useful event chain patterns.

The main contributions in this work are three-fold. (1) To support narrative understanding tasks with high-quality event chains, we develop a new event chain extraction method that is discourse-aware, and particularly, salience-aware. (2) We show the effectiveness of salience-based and discourse-aware filtering of event chains in improving narrative prediction, which leads to improvement on the Story Cloze Test. (3) We further demonstrate that the discourse patterns captured by the filtered event chains enhance language models on temporal understanding of events, leading to state-of-the-art performance on answering temporal ordering questions.

## 2 Event Chains and Narratives

Information extraction techniques have evolved to extract event mentions as well as their ordering information (Lin et al., 2020; Han et al., 2019; Wang et al., 2020), hence enabling the automated induction of raw event chains from text. Tools such as TEAR (Han et al., 2019) have been developed, which we can use to extract both the events and temporal ordering from text documents, like those shown in Figure 1. These sequences of events,

connected in a temporal order, form linear event chains, which can be viewed as a form of linear representation of the progress of described events in a narrative.

However, events that co-occur or are described together in an article are often not equally important. Some events are *salient* – they are relevant to the main topic of a text context, or play essential roles with regard to the central goal of an event chain (Liu et al., 2018; Chen et al., 2020). For example, from the paragraph in Figure 1, the events *detained*, *alerted* and *questioned* are salient components that constitute the progress of a described story. Other events may describe how the salient events came to be known or involve other events that happened alongside salient events but are not critical to understanding the core actions of the story. In Figure 1, *said* and *establish* are not salient. The goal of event salience detection systems (Liu et al., 2018; Jindal et al., 2020) is to identify events that are essential components of the narrative, which can help us filter out trivial or distracting event mentions from event chains.

From another related perspective, modeling discourse structures of the article is also helpful for understanding which parts of the text directly report on main events and which parts serve other supportive roles. Discourse profiling (Choubey et al., 2020) seeks to analyze the functional role of each sentence, which is different from the event-level prediction of salience. van Dijk (1988)’s theory of news discourse outlines an ontology of sentence types that are seen in most news articles. According to the theory, the first sentence in Figure 1 is a *Main Event* sentence, while the second is an *Evaluation* sentence. We may classify the members of van Dijk’s ontology into types that are core to understanding the event sequence of a story (e.g. *Main Event*) and types that are not (e.g. *Evaluation*), and further refine extracted event sequences.

Though salience and discourse structures represent different perspectives of narrative analysis, they can be concurrently modeled in an event chain extraction system, leading to more effective filtering of mined event chain representations.

## 3 Method

To obtain the interesting event chains from a document, we first use the TEAR tool by Han et al. (2019) to generate a temporal relation graph. Then we apply different levels of filtering on the ex-

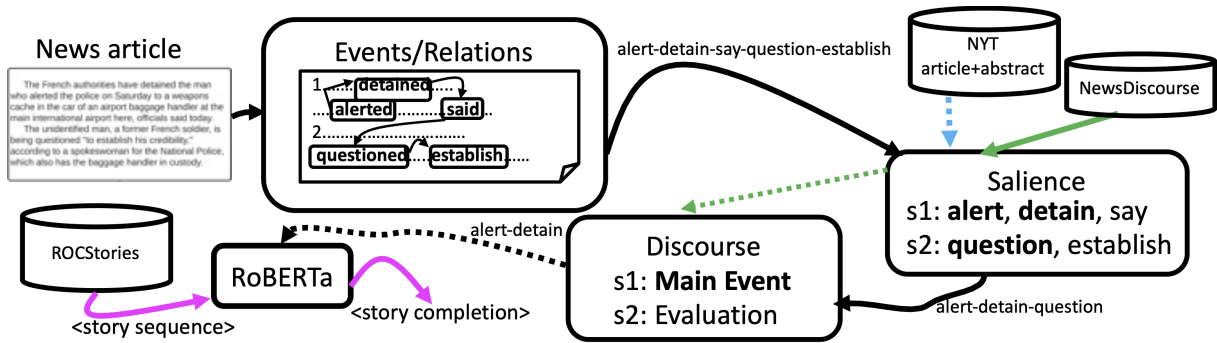


Figure 2: System diagram of our approach along with an example. Solid lines indicate inference and dotted lines indicate training. Colors separate different data streams. Events in temporal order are extracted from the news article from Figure 1. Saliency detection, trained on (article, abstract) pairs, filters out unimportant individual events from the example as well as labeled news discourse training data. Our saliency-aware discourse parsing model removes sentences that do not contribute to the direct story line. The important event chain is used to fine-tune a masked language model, which is used to predict story completion. A similar pipeline is used for the TORQUE task.

tracted events, which correspond to the nodes in the event graph structure. Finally, we extract linear chains of events from the filtered graph structure by following along the directed edges. In our evaluation, we will use these linear chains instead of raw text to pre-train language models such as RoBERTa (Liu et al., 2019). Our overall pipeline is shown in Figure 2. Details of each step in the pipeline are described as follows.

### 3.1 Event Chain Extraction

We adopt the joint structured event-relation extraction model from Han et al. (2019). It uses pre-trained BERT (Devlin et al., 2019) embedding vectors of the input text, which are further fed into an RNN-based scoring function for both event and relation extraction. During the end-to-end training, a MAP inference framework sets up a global objective function using local scoring functions to get the optimal assignment of event and relation labels. To ensure that we obtain globally coherent assignments, several logical constraints are specified including event-relation consistency and symmetry/transitivity consistency, so that the output event graph is logically consistent. This end-to-end model extracts both events and event relations simultaneously.

We restrict the event relation labels to *BEFORE* and *NONE*. We can regard each *BEFORE* relation as a directed edge in the temporal relation graph that points from the previous event to the next. Finally, to extract linear chains from this directed graph, we repeatedly choose the starting node in

a topological order, and start a walk along the directed edge to obtain a maximum chain.

### 3.2 Event Saliency Filtering

For the event saliency detection task, we follow the Kernel-based Centrality Estimation model by Liu et al. (2018). This model combines the various event saliency features such as frequency, sentence location and average cosine similarity with other events or entities, with Gaussian kernels that model event-event and entity-event interactions. For the final saliency score we apply an additional sigmoid function and use binary cross-entropy loss for training. In this way, we formulate the task as event-level binary classification.

We train the event saliency model on the New York Times (NYT) Annotated Corpus (Sandhaus, 2008), a collection of news articles with expert-written abstracts. During training, we use the frame-based event mention annotation by Liu et al. (2018), and for inference on new articles, we extract the event mentions using TEAR. Each event mention is labeled as salient if its lemma appears in the associated abstract. The trained saliency model can then assign a saliency score between 0 and 1 to each event extracted from a document, and we use a threshold of 0.5 to perform the final filtering of the events.

### 3.3 Discourse Saliency Filtering

For the discourse parsing model, we follow the hierarchical architecture by Choubey et al. (2020) to assign each sentence with one of the eight fine-




Figure 3: Architecture of salience-aware discourse parsing model. The addition of  $E_t$  encoding for each sentence, which is an average of the salience event encodings, provides additional event-level salience information not seen in the original discourse parsing model.

grained content labels<sup>2</sup> modified from van Dijk’s theory of news discourse (1988). The sentences that are labeled as either *Main Event*, *Consequence*, *Previous Event* or *Current Context* are considered relevant to the main sequence, and kept in the sentence-level filtering process.

One straightforward way to incorporate information from event salience in the filtering procedure is to first apply discourse filtering, and then keep only salient events from the discourse-filtered sentences, but this may propagate errors in the process and filter away too many events. We instead use event salience information to improve the performance on the discourse parsing task. As shown in Figure 3, we modify the hierarchical bi-LSTM model from the work by Choubey et al. (2020). The input is a document consisting of sentences  $s_1, s_2, \dots, s_n$ . Each sentence  $s_t$ , which is a sequence of tokens  $(w_{t1}, w_{t2}, \dots, w_{tk})$ , is first transformed to a sequence of ELMo (Peters et al., 2018) word representations, and then to hidden vectors  $(h_{t1}, h_{t2}, \dots, h_{tk})$  in the word-level bi-LSTM layer. Then, we concatenate the original intermediate sentence encoding  $S_t$  obtained from an soft-attention-weighted sum of the hidden vectors, with  $E_t$  which is an average over the hidden vectors of salient events within the sentence. These concatenated vectors are fed into the sentence-level bi-LSTM layer to generate sentence encodings

<sup>2</sup>Those include *Main Event* (M1), *Consequence* (M2), *Previous Event* (C1), *Current Context* (C2), *Historical Event* (D1), *Anecdotal Event* (D2), *Evaluation* (D3), and *Expectation* (D4).

$H_1, \dots, H_t, \dots, H_n$ , on top of which the final prediction layer and softmax are stacked, following the original model architecture by Choubey et al. (2020). Specifically, we compute the document encoding  $D$  as a soft-attention-weighted sum over the sentence encodings, and the final sentence representation of sentence  $t$  is the concatenation of  $H_t$  together with the element-wise product and difference between  $H_t$  and  $D$ . The final sentence representation is used in a two-layer feed-forward neural network to make final prediction of the sentence’s news discourse label.

We train the described model on the News-Discourse corpus with annotated content labels (Choubey et al., 2020), following their training setup. After we use our trained discourse parsing model on an input document, we filter the document down to only the salience-aware discourse-filtered sentences, which are sentences classified as one of *Main Event*, *Consequence*, *Previous Event* and *Current Context*, and keep only events from these sentences. We also try keeping only salient events from the filtered sentences, but it leads to worse performance as shown in our experiments.

### 3.4 Event Language Models

Once we obtain the linear event chains after performing the extraction and filtering from a text dataset, we treat the sequence of event mentions in each chain as a sequential context that would be used for training or fine-tuning a language model. For example, we can fine-tune a pre-trained Transformer language model based on the event chains (Section 4.1), or capture the sequences by training an RNN language model from scratch (Section 4.2). Once we obtain such a language model, we seek to use it to help with narrative prediction by predicting the continuation of an event chain. We can also leverage the fine-tuned model to support question answering regarding temporal orders of events.

## 4 Experiments

It is difficult to directly evaluate the quality of event chains in an intrinsic way. Some works on event schemas ask human annotators to rate qualities of generated chains (Balasubramanian et al., 2013; Weber et al., 2018a,b), which can be subjective. We instead turn to extrinsic evaluation tasks that depend on implicitly understanding typical sequences of meaningful events in order to be completed usefully.

TORQUE (Ning et al., 2020), short for *Temporal Ordering QUESTions*, is a machine reading comprehension benchmark that requests a model to answer questions regarding temporally specified events (e.g., “What happened before the snow started?”) in a reference article. We hypothesize that models trained on unfiltered event chains are less likely to focus on the relevant events requested by the question, but this seeks to be improved by our filtered event chains.

ROCStories (Mostafazadeh et al., 2016) is a narrative prediction dataset consisting of five-sentence short stories, where each sentence of a story contains a core event. A test set included with ROCStories contains two candidate endings to each partially complete story, where one of them is plausible. While prior work has successfully leveraged event chains to infer the endings (Chaturvedi et al., 2017), we believe a model trained on relevant chains of events should be able to better distinguish the relevant ending from the irrelevant one than a model trained on event chains that contain irrelevant events. Other works such as Sun et al. (2019) have directly tried to maximize performance on this corpus; we don’t seek to directly compete with that work here, but rather use ROCStories as a means of establishing that our isolation of important events at the event language model level does positively influence event ending prediction, indicating that the sequences we find do indeed contain more relevant events.

In this section we verify the value of our event-filtering model on these tasks. We also evaluate our proposed discourse parsing model to show the usefulness of combining salience information at multiple levels.

#### 4.1 Answering Temporal Ordering Questions

**Dataset** The TORQUE dataset has 3.2k news snippets and 21.2k user-provided questions. We follow the original data split given by Ning et al. (2020), using the training set for fine-tuning the language model for reading comprehension, and the dev set for evaluation. Each question asks about specific temporal relationships of the events in a news passage, and requests a sequence of event mentions from the passage that answer the question.

**Evaluation Details** We follow Ning et al. (2020) to fine-tune a RoBERTa-large (Liu et al., 2019) model on the training set of TORQUE, which has a perceptron layer that classifies whether each token

in the passage is in the answer or not. The input to the model is the question followed by the passage, separated by the separator token. We follow the same approach, but further fine-tune the model on inputs containing the extracted (and possibly filtered) *event chain*, rather than the entire passage. For training, we use the same batch size and learning rate as the baseline approach by Ning et al. (2020) in each experiment. We evaluate on the two standard metrics used in question answering, namely macro F1 and exact-match (EM), on the development set, comparing between models fine-tuned on filtered event chains and unfiltered chains, and, in Table 1, report the average performance over the 3 training runs started from different random seeds. As a comparison, we also fine-tune a GPT-2-based model (Radford et al., 2019), which leads to better results as shown in the rightmost column in Table 1.

**Results and Analysis** As we see in Table 1, we improve the performance on predictions of events over the baseline by fine-tuning on event chains constructed in various ways. Also, using filtered event chains gives better results than using unfiltered chains, regardless of the method of filtering. The model achieves the highest score when we use our salience-aware discourse filtering method on event chains, which shows the effectiveness of our approach of combining salience and discourse information. The last two rows of Table 1 show that keeping only salient events from the discourse-filtered sentences leads to worse results than keeping all events from the discourse-filtered sentences. We see the same trend in the last column, which suggests that the improvements from salience and discourse hold on GPT-2 as well as on RoBERTa fine-tuning. Comparing between the numbers from the two columns of “textual order”<sup>3</sup> and “TEAR order” in the table, we see the order of event chains does not seem to affect the performance much, probably because the number of events per document is not very large.

To further illustrate the effect of the various filtering methods on temporal understanding, we perform a breakdown of the questions in the dev set into different types, using prefix matching as shown in Figure 4. We define “standard” questions as those directly asking about Before, After or Co-occurring relations, without the additional

<sup>3</sup>TEAR is used to extract events but the temporal order relations are ignored.

Training Setting	Textual order		TEAR order		TEAR (GPT-2)	
	F1	EM	F1	EM	F1	EM
Baseline (Ning et al., 2020)	75.5	50.1	-	-	-	-
All events	75.9	50.2	75.8	50.2	78.9	53.5
Salient events	76.2	50.7	76.0	50.6	79.5	54.5
Discourse-filtered events	76.2	50.6	76.3	50.7	79.4	54.3
Saliency-aware discourse-filtered	76.8	51.0	76.9	51.2	<b>79.9</b>	<b>54.8</b>
Salient + Saliency-aware discourse-filtered	76.4	50.6	76.5	50.9	79.5	54.6

Table 1: Macro F1 and Exact-Match metrics (%) on TORQUE dev set under different training settings. The baseline setting corresponds to the original best performing setting of RoBERTa-large used in Ning et al. (2020) which does not use event chains, and the other settings refer to the different event-filtering methods. **Textual order** or **TEAR order** denotes if we construct the event chains in the order they appear in text or in the temporal order predicted by TEAR. The rightmost **TEAR (GPT-2)** column denotes the setting where we substitute RoBERTa-large with GPT-2.

Training Setting	Before	After	While
Baseline (Ning et al., 2020)	81.2	82.3	76.9
All events	82.9	83.7	77.5
Salient events	83.3	84.2	77.8
Discourse-filtered events	83.7	84.8	78.0
Saliency-aware discourse-filtered	<b>84.4</b>	<b>85.3</b>	<b>78.4</b>

Table 2: Macro F1 scores (%) on TORQUE dev set under different training settings, with a more detailed breakdown into “standard questions”. Event chains are in the temporal order predicted by TEAR, and other training settings are the same as in Table 1.

challenges from the fuzziness of time scopes or non-factual modality, as mentioned by Ning et al. (2020). The default questions ask about which events had already happened, were ongoing, or were still in the future, which we regard as not in the “standard” category. After examining the F1 scores on these “standard” categories in Table 2, we see that the model achieves a larger improvement with our filtering methods compared to on all questions, even though these are categories that already have a higher baseline performance.

## 4.2 Narrative Prediction

**Dataset** We evaluate the effectiveness of our extracted event chains on narrative prediction using the ROCStories dataset (Mostafazadeh et al., 2016). The training set contains 98,161 five-sentence stories, and the development and test sets each consist of 1,871 instances of four-sentence stories together with a correct and an incorrect ending. The goal is to predict the correct ending sentence given the two candidate choices. We use the development set but not the training set for supervised evaluation, following previous approaches (Chaturvedi et al., 2017; Srinivasan et al., 2018), as detailed below.

**Evaluation Details** Since each of the five sentences in a ROCStories article contains a core event, we convert the task to the prediction of the next event given the current sequence of events. We use the NYT corpus following Chaturvedi et al. (2017), and run our event extraction and filtering methods to obtain event chains from the news articles. We then break the obtained event chains into sequences of five events’ length, and train a simple bi-LSTM masked language model on these sequences. This seeks to compare with the best performing method using the same type of features (Chaturvedi et al., 2017). For evaluation on the Story Cloze Test, we choose the ending that contains the event with higher probability of occurring next given the four previous ones, according to the model. As a comparison, we also use the event chains to fine-tune a RoBERTa-large model. In this way we are performing an *unsupervised* evaluation in the sense that we are not learning directly from labeled sentences in ROCStories, and we are comparing the performance of models trained/fine-tuned on event chains obtained from different filtering methods. We also present results using the standard supervised setting (Chaturvedi et al., 2017; Srinivasan




Figure 4: Distribution of questions in the TORQUE dev set, showing only the most frequent prefixes related to Before, After or Co-occurring relations. For example, a total of 229 questions start with “What happened before”, which query the Before relation. 191+44 questions query the After relation, and 39+19 questions query the Co-occurring relation. We define these three categories of questions as “standard” type. This prefix matching is only a rough categorization of the question type.

et al., 2018), i.e., on top of the language model, we train a binary classifier using the correct and incorrect endings from the development set as positive and negative examples respectively.

**Results and Analysis** From Table 3, we can see the improvement of prediction accuracy when we feed filtered event chains to our model compared with using unfiltered chains. Our salience-aware discourse-filtered event chains, compared among all event chain-based training settings, provide the best results in both the unsupervised and supervised settings, regardless of the order of extracting the raw event chains. Under the unsupervised setting in which the model has not leveraged any supervision from the development set, and with event chains extracted in textual order, we see that event salience filtering and discourse filtering provide a 2.8% and 3.5% improvement on their own, respectively, and the largest improvement of 4.9% occurs when we use our proposed salience-aware discourse filtering. We see a similar trend under other evaluation settings, that combining event-level and discourse-level salience filtering leads to better performance, though the improvements are not as significant. This confirms that even without additional supervision from labeled data, we can utilize salience-aware and discourse-aware filtering to improve the relevance of event chains for better narrative prediction results. From the last row of Table 3, we see that our RoBERTa-based masked language model trained on the best produced event chains is able to further improve the performance in the supervised evaluation.

### 4.3 Salience-Aware Discourse Parsing

As a case study of our proposed discourse parsing model from Section 3.3, we also compare our performance on the discourse type classification task

with the baseline model from Choubey et al. (2020), as shown in Table 4. We see that our model surpasses the baseline model in both macro and micro F1 scores. Looking at the F1 score of each specific discourse type, we achieve the greatest improvement in the classification of type M1 (*Main Event*) and C2 (*Current Context*), with an increase of 5.5% and 6.3% respectively. This suggests that introducing salience awareness in the discourse parsing model indeed leads to better prediction accuracy, especially in categories that we expect to be most relevant to the topic of a document.

## 5 Related Work

We discuss three relevant research topics.

**Event Chains** Much research effort has been made to extract and process event chains. Chambers and Jurafsky (2008) pioneered the modern interest and modeled the co-occurrence of events in narrative chains based on their PMI. Radinsky et al. (2012) and Radinsky and Horvitz (2013) extended such unsupervised event chain modeling to cross-document scenarios and used the technology for news prediction and timeline construction. Berant et al. (2014) extracted relations among events and entities in biological processes to help solve a biological reading comprehension task using a structure matching method. More recently, Chen et al. (2020) attempted to infer the type of action and object associated with an event chain, which required recognition of the goal or intention extracted from the chain. Zhang et al. (2020) used a probabilistic graphical model to capture common patterns from analogous event chains, and induced new chains from those patterns. These works do not explore salience or discourse structures for event chains.

Training Setting	Textual order		TEAR order	
	unsupervised	supervised	unsupervised	supervised
Event-sequence SemLM (Chaturvedi et al., 2017)	-	71.6	-	-
All events	61.3	72.5	63.2	73.3
Salient events	64.1	73.4	63.6	73.8
Discourse-filtered events	64.8	73.6	63.8	74.1
Salience-aware discourse-filtered	66.2	74.2	65.8	74.5
Salient + Salience-aware discourse-filtered	65.0	73.9	64.4	74.2
Salience-aware discourse-filtered (RoBERTa)	<b>70.4</b>	<b>76.7</b>	<b>70.6</b>	<b>76.4</b>

Table 3: Accuracy (%) on ROCStories test set under different training settings. **Textual order** or **TEAR order** denotes if we construct the event chains from the NYT corpus in the order they appear in text or in the temporal order predicted by TEAR. The “unsupervised” mode only trains on sequences from the NYT documents, where the “supervised” mode also uses the development set of ROCStories. The baseline setting is a comparable logistic regression model trained on event-sequence modeling features only (Chaturvedi et al., 2017).

Models	M1	M2	C1	C2	D1	D2	D3	D4	Macro	Micro
Baseline	49.6	27.9	22.5	58.1	64.1	48.1	67.4	57.6	54.4	60.9
Ours	55.1	30.7	26.3	64.4	63.8	48.5	67.8	58.0	58.2	62.8

Table 4: Per-class and average F1 scores (%) of proposed salience-aware discourse parsing model, in the same evaluation setting as Choubey et al. (2020) (average of 10 runs with random seeds). Baseline model is the best performing one (**Document LSTM + Document Encoding**).

**Discourse and Salience** There have been recent interests in event salience identification. Choubey et al. (2018) studied how to find the dominant event(s) in a news article using mined event coreference relations. Liu et al. (2018) proposed a sequence tagging model for event salience detection, and also contributed a large-scale event salience corpus based on NYT. Jindal et al. (2020) tried to further capture representation of events and their interactions, and evaluated the modeling of event salience on extractive summarization.

Recent studies on discourse structures, especially for news articles, were built on Van Dijk’s theory (van Dijk, 1988). For example, Yarlott et al. (2018) annotated a dataset of discourse structures, which viewed paragraphs as units of annotations, and developed models to predict discourse labels. Choubey et al. (2020) instead created a sentence-level discourse structure corpus spanning different domains and sources, which was more helpful for a fine-grained identification of sentences relevant to the main topic. Insights in those studies constitute the two aspects of salience awareness that are characterized in our method, i.e., event-level and discourse-level salience.

**Event-Centric Language Models** Another line of research focuses on language modeling for capturing sequences of events. Many works in this line ex-

tract raw event chains and directly train a neural language model for narrative prediction (Chaturvedi et al., 2017; Peng et al., 2019) or script generation (Rudinger et al., 2015; Pichotta and Mooney, 2016; Weber et al., 2018b). Besides directly training, Zheng et al. (2020) proposed a unified fine-tuning architecture, where a masked language model was explicitly fine-tuned on the representations of event chains to model event elements. Li et al. (2020) proposed to learn to induce event schemas using an auto-regressive language model trained on salient paths on an event-event relation graph, which was also an attempt to reduce noise in constructing event schema. While that work uses a different kind of data structure, our works are in agreement on the importance of incorporating salience when fine-tuning an event-centric language model. Specifically, in comparison to prior works that used the language model for narrative prediction tasks (Chaturvedi et al., 2017; Peng et al., 2019), we show that salience-awareness is an important factor to tackle such tasks.

## 6 Conclusion

We propose an event chain extraction pipeline, which leverages both salience identification and discourse profiling to filter out distracting events. We demonstrate the effectiveness of the approach by



using the produced event chains to train/fine-tune language models, which leads to improved performance on temporal understanding of events and narrative prediction. Our case study on salience-aware discourse parsing shows the advantage of combining event-level and sentence-level salience information. We plan to use these event chain patterns on other narrative understanding and generation tasks, such as constrained story generation (Peng et al., 2018), event script generation (Zhang et al., 2020; Lyu et al., 2021), and implicit event prediction (Lin et al., 2021; Zhou et al., 2021).

## Ethical Considerations

This work does not present any direct societal consequence. The proposed method aims at providing high-quality extraction of event chains from documents with awareness of salience and discourse structures, but is only evaluated on English data and uses western notions of both salience and news discourse; event sequence extraction using data from other languages and cultures may not benefit from the methods shown here. The extracted event chain representations benefit narrative understanding and temporal understanding of events. Yet, real-world open source articles may include societal biases. Extracting event chains from articles with such biases may potentially propagate the bias into acquired knowledge representation. While not specifically addressed in this work, the ability to incorporate salience and discourse-awareness could be one way to mitigate bias.

## Acknowledgements

We appreciate the anonymous reviewers for their insightful comments and suggestions.

This research is based upon work supported by DARPA’s KAIROS program, Contract FA8750-19-2-0500. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

## References

Niranjan Balasubramanian, Stephen Soderland, Mausam, and Oren Etzioni. 2013. [Generating coherent event schemas at scale](#). In *Proceedings of the*

*2013 Conference on Empirical Methods in Natural Language Processing*, pages 1721–1731, Seattle, Washington, USA. Association for Computational Linguistics.

Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Abby Vander Linden, Brittany Harding, Brad Huang, Peter Clark, and Christopher D. Manning. 2014. [Modeling biological processes for reading comprehension](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1499–1510, Doha, Qatar. Association for Computational Linguistics.

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.

Snigdha Chaturvedi, Haoruo Peng, and Dan Roth. 2017. [Story comprehension for predicting what happens next](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1603–1614, Copenhagen, Denmark. Association for Computational Linguistics.

Muhao Chen, Hongming Zhang, Haoyu Wang, and Dan Roth. 2020. [What are you trying to do? semantic typing of event processes](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 531–542, Online. Association for Computational Linguistics.

Pengxiang Cheng and Katrin Erk. 2018. [Implicit argument prediction with event knowledge](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 831–840.

Prafulla Kumar Choubey, Aaron Lee, Ruihong Huang, and Lu Wang. 2020. [Discourse as a function of event: Profiling discourse structure in news articles around the main event](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5374–5386, Online. Association for Computational Linguistics.

Prafulla Kumar Choubey, Kaushik Raju, and Ruihong Huang. 2018. [Identifying the most dominant event in a news article by mining event coreference relations](#). 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 340–345, New Orleans, Louisiana. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*,

- pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rujun Han, Qiang Ning, and Nanyun Peng. 2019. [Joint event and temporal relation extraction with shared representations and structured prediction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 434–444, Hong Kong, China. Association for Computational Linguistics.
- Disha Jindal, Daniel Deutsch, and Dan Roth. 2020. [Is killed more significant than fled? a contextual model for salient event detection](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 114–124, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Manling Li, Qi Zeng, Ying Lin, Kyunghyun Cho, Heng Ji, Jonathan May, Nathanael Chambers, and Clare Voss. 2020. [Connecting the dots: Event graph schema induction with path language modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 684–695, Online. Association for Computational Linguistics.
- Shih-Ting Lin, Nathanael Chambers, and Greg Durrett. 2021. [Conditional generation of temporally-ordered event sequences](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7142–7157, Online. Association for Computational Linguistics.
- Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. [A joint neural model for information extraction with global features](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7999–8009, Online. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv:1907.11692*.
- Zhengzhong Liu, Chenyan Xiong, Teruko Mitamura, and Eduard Hovy. 2018. [Automatic event salience identification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1226–1236, Brussels, Belgium. Association for Computational Linguistics.
- Qing Lyu, Li Zhang, and Chris Callison-Burch. 2021. [Goal-oriented script construction](#). In *Proceedings of the 14th International Conference on Natural Language Generation*.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. [TORQUE: A reading comprehension dataset of temporal ordering questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1158–1172, Online. Association for Computational Linguistics.
- Simon Ostermann, Michael Roth, and Manfred Pinkal. 2019. [MCScript2.0: A machine comprehension corpus focused on script events and participants](#). In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (\*SEM 2019)*, pages 103–117, Minneapolis, Minnesota. Association for Computational Linguistics.
- Haoruo Peng, Qiang Ning, and Dan Roth. 2019. [KnowSemLM: A knowledge infused semantic language model](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 550–562, Hong Kong, China. Association for Computational Linguistics.
- Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. 2018. [Towards controllable story generation](#). In *Proceedings of the First Workshop on Storytelling*, pages 43–49, New Orleans, Louisiana. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Karl Pichotta and Raymond J. Mooney. 2016. Learning statistical scripts with LSTM recurrent neural networks. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI’16*, page 2800–2806. AAAI Press.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Kira Radinsky, Sagie Davidovich, and Shaul Markovitch. 2012. [Learning causality for news events prediction](#). In *Proceedings of the 21st International Conference on World Wide Web, WWW ’12*, page 909–918, New York, NY, USA. Association for Computing Machinery.

- Kira Radinsky and Eric Horvitz. 2013. [Mining the web to predict future events](#). In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining, WSDM '13*, page 255–264, New York, NY, USA. Association for Computing Machinery.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. [CoQA: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Rachel Rudinger, Pushpendre Rastogi, Francis Ferraro, and Benjamin Van Durme. 2015. [Script induction as language modeling](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1681–1686, Lisbon, Portugal. Association for Computational Linguistics.
- Evan Sandhaus. 2008. *The New York Times Annotated Corpus*. LDC corpora. Linguistic Data Consortium.
- Siddarth Srinivasan, Richa Arora, and Mark Riedl. 2018. [A simple and effective approach to the story cloze test](#). 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 92–96, New Orleans, Louisiana. Association for Computational Linguistics.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2019. [Improving machine reading comprehension with general reading strategies](#). 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 2633–2643, Minneapolis, Minnesota. Association for Computational Linguistics.
- T.A. van Dijk. 1988. *News as Discourse*. Communication (Hillsdale, N.J.). L. Erlbaum Associates.
- Haoyu Wang, Muhao Chen, Hongming Zhang, and Dan Roth. 2020. [Joint constrained learning for event-event relation extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 696–706, Online. Association for Computational Linguistics.
- Noah Weber, Niranjan Balasubramanian, and Nathanael Chambers. 2018a. [Event representations with tensor-based compositions](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Noah Weber, Leena Shekhar, Niranjan Balasubramanian, and Nathanael Chambers. 2018b. [Hierarchical quantized representations for script generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3783–3792, Brussels, Belgium. Association for Computational Linguistics.
- Margaret Wilson. 2002. [Six views of embodied cognition](#). *Psychonomic bulletin & review*, 9(4):625–636.
- Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. [Plan-and-write: Towards better automatic storytelling](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7378–7385.
- W. Victor Yarlott, Cristina Cornelio, Tian Gao, and Mark Finlayson. 2018. [Identifying the discourse function of news article paragraphs](#). In *Proceedings of the Workshop Events and Stories in the News 2018*, pages 25–33, Santa Fe, New Mexico, U.S.A. Association for Computational Linguistics.
- Hongming Zhang, Muhao Chen, Haoyu Wang, Yangqiu Song, and Dan Roth. 2020. [Analogous process structure induction for sub-event sequence prediction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1541–1550.
- Jianming Zheng, Fei Cai, and Honghui Chen. 2020. [Incorporating scenario knowledge into a unified fine-tuning architecture for event representation](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, page 249–258, New York, NY, USA. Association for Computing Machinery.
- Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2021. [Temporal reasoning on implicit events from distant supervision](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1361–1371, Online. Association for Computational Linguistics.