

Uncovering Surprising Event Boundaries in Narratives

Anonymous ACL submission

Abstract

When reading stories, people can naturally identify sentences in which a new event starts, i.e., *event boundaries*, using their knowledge of how events typically unfold, but a computational model to detect event boundaries is not yet available. We characterize and detect sentences with expected or surprising event boundaries in an annotated corpus of short diary-like stories, using a model that combines commonsense knowledge and narrative flow features with a RoBERTa classifier. Our results show that, while commonsense and narrative features can help improve performance overall, detecting event boundaries that are more subjective remains challenging for our model. We also find that sentences marking surprising event boundaries are less likely to be causally related to the preceding sentence, but are more likely to express emotional reactions of story characters, compared to sentences with no event boundary.

1 Introduction

When people read stories, they can easily detect the start of new events through changes in circumstances or in narrative development, i.e., *event boundaries* (Zacks et al., 2007; Bruni et al., 2014; Foster and Keane, 2015; Jafarpour et al., 2019b). These event boundaries can be expected or surprising. For example, in the story in Figure 1 based on crowdsourced annotation, “getting along with a dog who does not generally like new people” marks a *surprising* new event, while “their playing fetch together for a long time” is an *expected* new event.

We aim to study whether machines can detect these surprising or expected event boundaries, using commonsense knowledge and narrative flow features. Characterizing features that are informative in detecting event boundaries can help determine how humans apply expectations on event relationships (Schank and Abelson, 1977; Kurby and Zacks, 2009; Radvansky et al., 2014; Ünal



Figure 1: Example story with sentences that contain either a surprising event boundary, no event boundary or an expected event boundary respectively. The annotations of reader perception are from the Hippocampus dataset (Sap et al., 2022).

et al., 2019; Zacks, 2020). Furthermore, detection of sentences with event boundaries can also be useful when generating engaging stories with a good amount of surprises. (Yao et al., 2019; Rashkin et al., 2020; Ghazarian et al., 2021).

To differentiate sentences with surprising event boundaries, expected event boundaries, and no event boundaries, we train a classifier using 3925 story sentences with human annotation of event boundaries from diary-like stories about people’s everyday lives (Sap et al., 2022). We extract various commonsense and narrative features on relationships between sentences of a story, which can predict the type of event boundaries. Commonsense features include the likelihood that adjacent sentences are linked by commonsense relations from the knowledge graphs Atomic (Sap et al., 2019a) and Glucose (Mostafazadeh et al., 2020). Narrative features include Realis (Sims et al., 2019) that identifies the number of event-related words in a sentence, Sequentiality (Radford et al., 2019; Sap et al., 2022) based on the probability of generating a sentence with varying context and SimGen (Rosset, 2020), which measures the similarity between a sentence and the sentence that is most likely to

067 be generated given the previous sentence. We then
068 combine the prediction based on these features with
069 the prediction from a RoBERTa classifier (Liu et al.,
070 2019), to form overall predictions.

071 We evaluate the performance of the classification
072 model by measuring F1 of the predictions and com-
073 pare various configurations of the model to a base-
074 line RoBERTa model. We find that integrating nar-
075 rative and commonsense features with RoBERTa
076 leads to a significant improvement (+2.2% F1) over
077 a simple RoBERTa classifier. There are also in-
078 dividual differences on the subjective judgment
079 of which sentences contain a surprising or an ex-
080 pected event boundary, that is reflected in the detec-
081 tion model’s performance. The performance of our
082 model increases with increasing agreement across
083 the human annotators. Additionally, by interpreting
084 the trained parameters of our model, we find that
085 the absence of causal links between sentences is a
086 strong predictor of surprising event boundaries.

087 To further analyze how surprising event bound-
088 aries relate to deviation from commonsense un-
089 derstanding, we compare the performance of the
090 classification model on the related task of ROC
091 Story Cloze Test (Mostafazadeh et al., 2016). This
092 task concerns whether the ending sentence of a
093 story follows/violates commonsense based on ear-
094 lier sentences, which can be linked to whether
095 sentences are expected or surprising. Our model
096 performs significantly higher on the ROC Story
097 Cloze Test (87.9% F1 vs 78.0% F1 on our task),
098 showing that surprising event boundaries go be-
099 yond merely violating commonsense and therefore
100 can be seen as more challenging to detect. Together,
101 our results suggests that while detecting surprising
102 event boundaries remains a challenging task for
103 machines, a promising direction lies in utilizing
104 commonsense knowledge and narrative features to
105 augment language models.

106 2 Event Boundary Detection Task

107 Events have been widely studied in Natural Lan-
108 guage Processing. They have often been repre-
109 sented in highly structured formats with word-
110 specific triggers and arguments (Walker et al., 2006;
111 Li et al., 2013; Chen et al., 2017; Sims et al., 2019;
112 Mostafazadeh et al., 2020; Ahmad et al., 2021)
113 or as Subject-Verb-Object-style (SVO) tuples ex-
114 tracted from syntactic parses (Chambers and Ju-
115 rafsky, 2008; Martin et al., 2018; Rashkin et al.,
116 2018; Sap et al., 2019a). In narratives, events

117 are represented as a continuous flow with multiple
118 boundaries marking new events (Zacks et al., 2007;
119 Graesser et al., 1981; Kurby and Zacks, 2008; Za-
120 cks, 2020); however, we lack a model to detect the
121 boundary events that mark the meaningful segmen-
122 tation of a continuous story into discrete events.

123 In this work, we study stories from a cognitive
124 angle to detect event boundaries. Such event bound-
125 aries relate to our narrative schema understanding
126 (Schank and Abelson, 1977; Chambers and Juraf-
127 sky, 2008; Ryan, 2010), commonsense knowledge
128 (Sap et al., 2019a; Mostafazadeh et al., 2020) and
129 world knowledge (Nematzadeh et al., 2018; Bisk
130 et al., 2020). Event boundaries can be surprising or
131 expected based on the knowledge of how a flow of
132 events should unfold. For example, events can be
133 surprising when they deviate from commonsense in
134 terms of what people would predict (e.g., if some-
135 one won something, they should not be sad; Sap
136 et al., 2019a). Surprising events can also be low
137 likelihood events (Foster and Keane, 2015) such as
138 seeing someone wear shorts outside in winter, or
139 due to a rapid shift in emotional valence between
140 events (Wilson and Gilbert, 2008) such as seeing
141 a protagonist being defeated. Importantly, there
142 are individual differences in how humans segment
143 narratives into events (Jafarpour et al., 2019a).

144 We tackle event boundary detection as a three-
145 way classification task that involves distinguishing
146 surprising but plausible event boundaries in story
147 sentences from expected event boundaries and no
148 event boundaries. To mirror how humans read sto-
149 ries, we predict the event boundary label for a sen-
150 tence using all of its preceding sentences in the
151 story, as well as the general story topic as context.
152 *Surprising* event boundaries are novel events that
153 are unexpected given their context, such as a dog
154 getting along with someone despite not typically
155 liking new people. *Expected* event boundaries are
156 novel events that are not surprising, such as a per-
157 son playing a new game with a dog for a long time
158 given that they like each other. In contrast, sen-
159 tences with *no event* boundary typically continue
160 or elaborate on the preceding event, such as a per-
161 son liking a dog given that they get along with the
162 dog (Figure 1).

163 3 Event-annotated Data

164 We use the event-annotated sentences from stories
165 in the Hippocorpus dataset to study event bound-
166 aries. This dataset contains 240 diary-like sto-

Majority label	#Samples (%)	% majority agreement (std)
No event	2255 (57.5)	68.1 (13.9)
Expected	650 (16.6)	58.8 (10.6)
Surprising	509 (13.0)	61.7 (11.9)
Tied	511 (13.0)	41.1 (5.7)
Total	3925 (100)	62.2 (15.2)

Table 1: Descriptive Statistics for Event-Annotated sentences. Majority label refers to the most common annotation of a sample from 8 independent annotators. If there is a tie between 2 labels, it is categorized as tied. Majority agreement is the proportion of sample annotations for the majority label.

167 ries about everyday life experiences, which annotated at the sentence level (Sap et al., 2022). Stories were inspected for the absence of offensive or person-identifying content. For the annotation, eight crowdworkers were shown a story sentence by sentence and were asked to mark whether each sentence contained a new surprising or expected event boundary, or no event boundary at all, based on their subjective judgment (Sap et al., 2022). Summarized in Table 1, based on the majoritarian vote, most sentences (57.5%) contain no event boundaries while 16.6% and 13.0% of sentences contains expected and surprising event boundaries, respectively.

168 Due to the inherent subjectivity of the task, aggregating labels into a majority label yields low agreement (e.g., 61.7% for surprising event boundaries; Table 1). Therefore, at training time, we use the proportion of annotations for each event boundary type as the label instead of the majority vote, because such distributional information is a better reflection of the inherent disagreement among human judgements (Pavlick and Kwiatkowski, 2019). At test time, we use the majority vote as a gold label, since measuring performance on distribution modelling is less intuitive to interpret, and subsequently break down performance by agreement level to take disagreements into account.

4 Event Boundary Detection Model

196 We first describe informative commonsense and narrative features that we extract for the event boundary detection model. Then, we describe how we integrate these features with a RoBERTa classifier in our model before detailing our experimental setup. Figure 2 depicts an overview of our model.

4.1 Features

202 We select a collection of commonsense features (Atomic and Glucose relations) and narrative flow features (Realis, Sequentiality and SimGen). A model is trained separately from our main model for Atomic relations, Glucose relations and Realis while models for Sequentiality and SimGen are used without further training. Features of story sentences are extracted as input into the main model. Because language modelling alone might not be sufficient to learn such features (Gordon and Van Durme, 2013; Sap et al., 2019a), we provide the extracted features to the model instead of relying on the language models to learn them implicitly.

217 **Atomic relations** are event relations from a social commonsense knowledge graph containing numerous events that can be related to one another (Sap et al., 2019a). The event relations in this graph consists of:

- 222 Emotional **Reaction**,
- 223 The **Effect** of an event,
- 224 **Want** to do after the event,
- 225 What **Needs** to be done before an event,
- 226 The **Intention** to do a certain event,
- 227 What **Attributes** an event expresses.

228 When an event affects the subject, the feature name is preceded by an \times , while if it affects others, it has an \circ . For example, an \times Want of a sentence *PersonX pays PersonY a compliment* is that *PersonX will want to chat with PersonY*, and an \circ Want is that *PersonY will compliment PersonX back*. We use Atomic relations because surprising event boundaries can involve breaches of commonsense understanding (Bosselut et al., 2019; Sap et al., 2019a; Mostafazadeh et al., 2020; Gabriel et al., 2021). Furthermore, some Atomic relations (\times React and \circ React) concern emotional affect and therefore can be used to capture changes in emotional valence, which can cause events to be seen as surprising (Wilson and Gilbert, 2008).

243 We train an Atomic relation classifier using a RoBERTa-base model (Liu et al., 2019) to classify event-pairs into one of the nine possible relationship labels as well as a None label (to introduce negative samples). We achieved a validation F1 of 77.15%, which is high for a 10-way classification task. We describe training and other experimental details in the Appendix. When making inferences on the event-annotated dataset, we predict the likelihood that a preceding sentence in a story will be

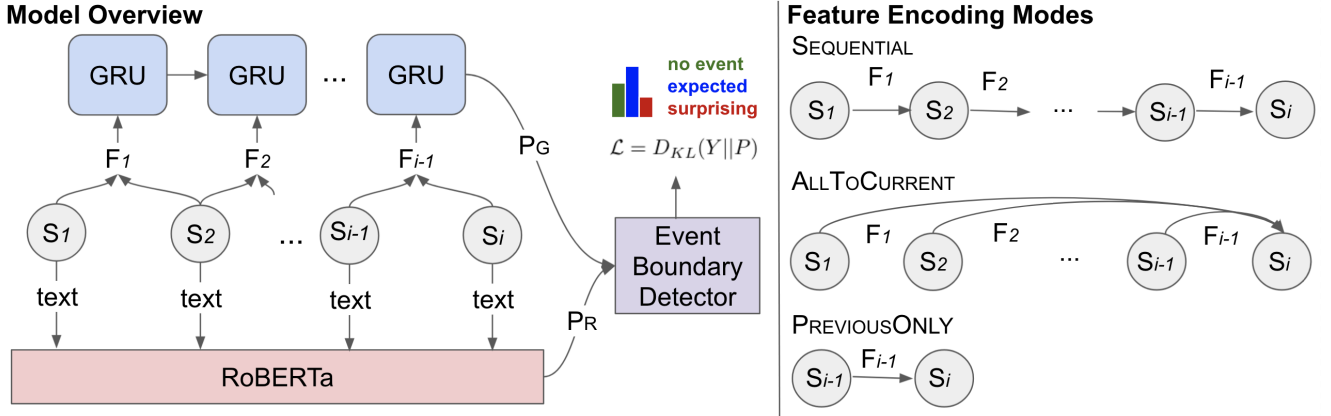


Figure 2: (Left) Our model involves a **GRU** to combine features from sentence pairs with three feature encoding modes, **RoBERTa** to consider story sentences and **Event Boundary Detector** to combine predictions made by the two components. S_n and F_n refer to sentence n and features n respectively, while P_G and P_R are predictions made by the GRU and RoBERTa. The output is a probability distribution over no event boundary, expected event boundary and surprising event boundary, which is used to update model parameters together with the label using the **Kullback-Leibler Divergence** loss function. (Right) **Features** (Atomic, Glucose, Realis, Sequentiality and SimGen) can be extracted as input into the GRU in three feature encoding modes: SEQUENTIAL (shown in Model Overview), ALLTOCURRENT and PREVIOUSONLY.

related to the current sentence via each of the nine relationship label. Because Atomic relations are directed relations (e.g., *I ate some cake* xEffect *I am full* is different from *I am full* xEffect *I ate some cake*), we also made the reverse inference in case commonsense relations between sentences exist in the reverse direction. Together, 9 forward atomic relation features and 9 reverse features (marked with ‘-r’) are used.

Glucose relations are event relations from another commonsense knowledge dataset containing relations between event-pairs in 10 dimensions (Mostafazadeh et al., 2020). Glucose relation features are used to complement Atomic relation features in its coverage of commonsense relations. Dim-1 to 5 are described below while Dim-6 to 10 are the reverse/passive form of Dim-1 to 5 respectively.

- Dim-1: **Event** that causes/enables
- Dim-2: **Emotion/human drive** that motivates
- Dim-3: **Change in location** that enables
- Dim-4: **State of possession** that enables
- Dim-5: **Other attribute** that enables

Glucose relation classifier was trained on a RoBERTa-base model to classify event-pairs from its annotated dataset into one of ten possible relation labels as well as a None label. We used the *specific* version of Glucose events represented in natural language. As a result, we achieved a validation F1 of 80.94%. Training and other experimental

details are in the Appendix. During inference on the Event-annotated dataset, we predict and use as features the likelihood that the current sentence will be related to a preceding sentence via each relation label.

Realis events are words that serve as triggers (i.e., head words) for structured event representations (Sims et al., 2019). Realis event words denote concrete events that actually happened, meaning that a higher number of Realis event words suggests greater likelihood of the sentence containing a new event boundary (expected or surprising). We trained a BERT-base model (Devlin et al., 2019) on an annotated corpus of literary novel extracts (Sims et al., 2019). We achieved a validation F1 of 81.85%, inspired by and on par with Sap et al. (2020). Then, we use the trained model to make inference on story sentences in the Event-annotated dataset. Finally, we used the number of Realis words in each sentence as a feature. Training and other experimental details are in the Appendix.

Sequentiality is a measure of the difference in conditional negative log-likelihood of generating a sentence given the previous sentence or otherwise (Sap et al., 2020, 2022). Sequentiality can be a predictor for unlikely events, which can cause surprise (Foster and Keane, 2015). We use GPT-2 (Radford et al., 2019) to measure this negative log-likelihood since it is a Left-to-Right model, which

312 matches the order in which annotators were shown 359
 313 sentences in a story. NLL of each sentence was 360
 314 obtained in two different contexts. `NLL_topic` 361
 315 is based on the sentence alone with only the topic 362
 316 as prior context, while `NLL_topic+prev` uses 363
 317 the previous sentence as additional context to 364
 318 study the link between adjacent sentences. Finally, 365
 319 `Sequentiality` is obtained by taking their dif- 366
 320 ference. Experimental details are in the Appendix. 367

$$321 \quad NLL_{topic} = -\frac{1}{|s_i|} \log p_{LM}(s_i | Topic)$$

$$322 \quad NLL_{topic+prev} = -\frac{1}{|s_i|} \log p_{LM}(s_i | Topic, s_{i-1})$$

323 **SimGen** is computed as the cosine similarity be- 374
 324 tween each sentence and the most likely gener- 375
 325 ated sentence given the previous sentence, under 376
 326 a large Left-to-Right language model (specifically, 377
 327 Turing-NLG; Rosset, 2020). Then, we separately 378
 328 converted the original sentence and generated sen- 379
 329 tence into sentence embeddings using a pre-trained 380
 330 MPnet-base model (Song et al., 2020). Finally, the 381
 331 generated embeddings and the original embeddings 382
 332 are compared for cosine similarity, which is used as 383
 333 a feature. Experimental details are in the Appendix. 384

334 4.2 Model Architecture 385

335 We propose a model to integrate feature-based pre- 386
 336 diction with language-based prediction of event 387
 337 boundaries, illustrated in Figure 2 (left). The pre- 388
 338 dictions are independently made with extracted fea- 389
 339 tures using a gated recurrent unit (GRU) and with 390
 340 language (i.e., story sentences) using RoBERTa. 391
 341 Then these predictions are combined into a final 392
 342 predicted distribution for the three types of event 393
 343 boundaries. Our model is then trained using the 394
 344 Kullback-Leibler Divergence loss. 395

345 **GRU** is used to combine features relating the cur- 396
 346 rent sentence i to prior sentences in a story. It se- 397
 347 quentially considers information concerning prior 398
 348 sentences, which mimics the annotator’s procedure 399
 349 of identifying event boundaries as they read one 400
 350 sentence at the time. As seen in Figure 2 (right), 401
 351 we use three feature encoding modes to determine 402
 352 the features that are used as input into the GRU, as 403
 353 inspired by literature on event segmentation (Petti- 404
 354 john and Radvansky, 2016; Baldassano et al., 2018; 405
 355 Zacks, 2020). These three modes represents differ- 406
 356 ent ways of facilitating information flow between 407
 357 sentences, which can have distinct effects on iden-
 358 tifying event boundaries.

The first mode, `SEQUENTIAL`, encodes features 359
 from all previous sentences in the story in a re- 360
 current way (1 to 2, 2 to 3 ... $i - 1$ to i) up until 361
 the current sentence i . The second mode, `ALL-` 362
`TOCURRENT`, uses features from each of the previ- 363
 ous sentences to the current sentence i (1 to i , 2 to 364
 $i ... i - 1$ to i). The third mode, `PREVIOUSONLY`, 365
 ($i - 1$ to i) only feeds into the GRU the features 366
 relating to the previous sentence. For all modes, the 367
 dimension of each time step input is K_G , represent- 368
 ing the total number of distinct features. We then 369
 project the final output of the GRU, $h_G \in \mathbb{R}^{K_G}$, 370
 into a 3-dimensional vector space representing the 371
 unnormalized probability distribution over event 372
 boundary types. 373

RoBERTa is used to make predictions based on 374
 text in story sentences. We use all story sentences 375
 up to sentence i inclusive. We then project the 376
 hidden state of the first token, $h_R \in \mathbb{R}^{K_R}$, into a 3- 377
 dimensional space representing the unnormalized 378
 probability distribution over event boundary types. 379

Combining predictions We combine predic- 380
 tions made by the GRU (P_G) and RoBERTa (P_R) 381
 by concatenating their predictions and multiplying 382
 it with a linear classifier of size (6, 3) to output 383
 logits of size (3). The logits are then normalized 384
 using Softmax to give a distribution of the three 385
 types of event boundaries (P). The weights of the 386
 linear classifier are initialized by concatenating two 387
 identity matrix of size 3 (\mathbf{I}_3), which serves to per- 388
 form elementwise addition between the predictions 389
 of the GRU and RoBERTa at early stages of the 390
 training process. 391

$$392 \quad W := [\mathbf{I}_3; \mathbf{I}_3] \quad (1)$$

$$393 \quad P := \text{Softmax}(W([P_G; P_R])) \quad (2) \quad 394$$

Loss function We use the Kullback-Leibler Di- 395
 vergence loss function to train the model. We use 396
 it over the standard Cross Entropy loss function 397
 because our training targets are in the form: propor- 398
 tion of annotations for each type of event boundary 399
 (e.g., 0.75, 0.125, 0.125 for no event, expected 400
 and surprising respectively). Including such dis- 401
 tributional information in our training targets over 402
 using the majority annotation only can reflect the 403
 inherent disagreement among human judgements 404
 (Pavlick and Kwiatkowski, 2019), which is impor- 405
 tant to capture for event boundaries given that they 406
 are subjective judgements. 407

4.3 Experimental setup

We seek to predict the event-boundary annotation for each Hippocampus story sentence, using preceding sentences in the story as context, as shown in Figure 2. Additional training and experimental details are available in the Appendix.

K-fold Cross-validation Because of the limited size of the dataset ($n=3925$), we split the dataset in k -folds ($k=10$), using one fold ($n=392$) for validation and nine other folds combined for training. From each of the 10 models, we obtained the prediction for the validation set. Together, the validation sets for the 10 models combine to form predictions for the entire dataset, which we use to conduct significance testing in order to compare the performance of models.

GRU was accessed from PyTorch, with K_G set to 33 and a hidden dimension of 33.

RoBERTa RoBERTa-base-uncased was used, accessed from HuggingFace Transformers library (with 12-layer, 768-hidden (K_R), 12-heads, 110M parameters, 0.1 dropout). When more than 10 prior sentences are available in a story, we use only the most recent 10 sentences due to RoBERTa input sequence length limitations.

Evaluation Metrics While capturing distributional information of subjective judgement labels (Pavlick and Kwiatkowski, 2019) is important for training, it can also be difficult to interpret for evaluation. Therefore, we decided to predict for the most likely label during evaluation and compare it against the majority label for each sample. Some samples do not have a single majority label (e.g., equal number of expected and surprising annotations) and these samples were excluded. We use micro-averaged F1 as the metric.

5 Results and Discussion

We first quantify the performance of our model in detecting event boundaries, using a coarse-grained performance measure on F1 with respect to majority vote. Then, we investigate how the performance varies based on annotation subjectivity. Finally, we inspect the model parameters to identify commonsense and narrative features that are most informative in detecting event boundaries.

Improving prediction of event boundaries As seen in Table 2, RoBERTa alone performs fairly

	overall F1	no event F1	expected F1	surprising F1
Event Detector (w RoBERTa)				
- PREVIOUSONLY*	78.0	87.2	60.0	59.7
- SEQUENTIAL	77.3	86.6	57.5	60.5
- ALLTOCURRENT	76.9	86.3	57.5	59.7
RoBERTa	75.8	86.2	55.8	54.3
Event Detector (w/o RoBERTa)				
- ALLTOCURRENT	63.9	81.8	32.3	24.8
- SEQUENTIAL	63.8	82.1	34.6	19.5
- PREVIOUSONLY	63.4	81.8	31.8	21.2

Table 2: Event detection task: Performance of Event Detector compared to baseline model. *: significant difference from RoBERTa based on McNemar’s test ($p < 0.05$)

well in predicting event boundaries (F1 = 75.8%, within 2.2% F1 of our best performing model), but can be further supported by our commonsense and narrative features to improve its performance. In contrast, the commonsense and narrative features alone do not perform as well.¹ Overall, our best performing set-up is the Event Detector (PREVIOUSONLY) with F1 = 78.0%, which is significantly different from RoBERTa alone based on McNemar’s test ($p < 0.05$).² Its overall strong performance is largely contributed by its strong performance in detecting no event boundaries and expected event boundaries. F1 for no event boundary is higher than for both surprising and expected event boundaries, likely because there are more sentences with no event boundaries as seen in Table 1. The PREVIOUSONLY configuration performs best for no event boundaries and expected event boundaries likely because determining whether the current sentence continues an expected event (or not) requires retaining the latest information in working memory (Jafarpour et al., 2019a). However, the SEQUENTIAL configuration seems to perform the best in predicting surprising event boundaries. Compared to no/expected event boundaries, we hypothesize that predicting surprising event boundaries requires taking into account how the story developed prior to the previous sentence in setting up the context for the current sentence. This

¹We also increased learning rate to $1e-3$ for better performance given the absence of RoBERTa predictions in this ablation set-up

²McNemar’s test is used to determine whether samples that have been predicted accurately (or not) by one model overlap with those that have predicted accurately (or not) by another model

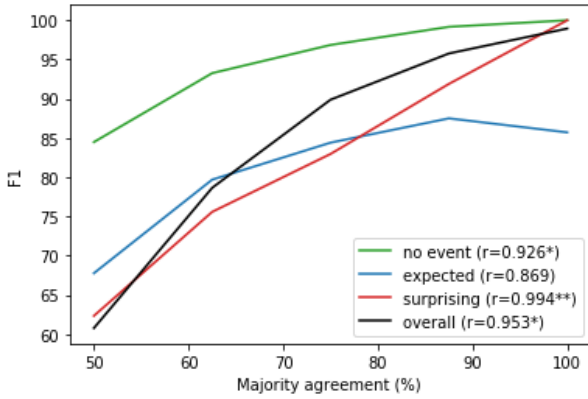


Figure 3: F1 by Event Detector (PREVIOUSONLY) against majority agreement, on all 10 folds. * means that Pearson’s r is significant at $p < 0.05$ and ** at $p < 0.001$.




Figure 4: Feature weights towards each label in GRU component of Event Detector (PREVIOUSONLY)

484 finding echoes results by Townsend (2018) that
 485 showed that surprising sentences take long time to
 486 read because it requires changing our mental model
 487 formed from previous sentences.

488 **F1 varies with majority agreement** Since the
 489 annotations were subjective and did not always
 490 agree, we further examine our best model’s per-
 491 formance (PREVIOUSONLY) with respect to an-
 492 notation agreement. As shown in Figure 3, F1
 493 increases with majority label agreement (Pearson’s
 494 $r = 0.953$, $p < 0.05$). Such positive correlations
 495 are observed across all event boundary labels (Pear-
 496 son’s $r = 0.869-0.994$) and is especially strong for
 497 surprising event boundaries (Pearson’s $r = 0.994$,
 498 $p < 0.001$). This means that most errors are made
 499 on samples that have low agreement among anno-
 500 tators. For example to show this contrast, after
 501 “*She and I are very close so it was great to see her*
 502 *marrying someone she loves,*” 7 out of 8 annotators
 503 indicated that “*The most memorable moment was*
 504 *when I spilled champagne on my dress before the*
 505 *wedding*” was surprising. On the other hand, after
 506 “*It was a hot day in July that our community de-*
 507 *ecided to paint a mural on an intersection for public*
 508 *art,*” only 4 out of 8 annotators indicated that “*I had*
 509 *decided to volunteer to help paint.*” was surprising.
 510 The results suggest that our model performance
 511 reflects the variability and agreements in humans
 512 annotations of event boundaries. We hypothesize
 513 that the event boundaries with more agreement are
 514 based on features that are shared across the an-
 515 notators, such as commonsense knowledge; there-
 516 fore, the model performs well in detecting those.
 517 Whereas, our model struggles with detecting event

boundaries that are more subjective.

Predictive features By integrating a separate
 feature-based classifier, the Event Boundary Detec-
 tor model allows us to examine the model paramet-
 ers and determine features that are associated with
 surprising, expected or no event boundaries. First,
 we take the average of the GRU classifier weights
 for each of the 10 cross-validated models. Then,
 we plot these weights for each label in Figure 4,
 and summarize the findings below.

Features that relate to commonsense relations:
 oEffect, xEffect and Glucose Dim-6
 (caused by) are most predictive of expected event
 boundaries. This can indicate that events that are an
 effect of/caused by a prior event can be expected by
 annotators, as also noted by Graesser et al. (1981).
 An example of an expected event boundary is “*I*
told her we could go for coffee sometime.”, as an
 effect of “*We had a good time together.*” xNeed is
 least indicative of surprising event boundaries. This
 is likely because xNeed refers to what the subject
 need to do before an activity, which is procedural
 and unlikely to cause surprise. An example is “*I*
was grocery shopping a few weeks ago.” which is
 needed before “*I had purchased my items and was*
leaving the store.”

Features that explain unlikely events Realis
 is highest for surprising event boundaries, suggest-
 ing that surprising event boundaries tend to con-
 tain the most concrete event-words. Surprising
 event boundaries also have the highest likelihood
 when conditioned on the story topic (NLL_topic)
 while expected events are highest when condi-
 tioned based on the topic and the previous sentence

(NLL_topic+prev). This suggests that surprising events are often inline with the story topic but not with the previous sentence. Therefore, the low likelihood of transitioning between the previous and current sentence is a strong predictor of surprising event boundaries, in line with findings by Foster and Keane (2015) on how the difficulty of linking two adjacent events is an important factor in causing surprise.

Features that explain changes in emotional valence Compared to sentences that contain no event boundaries, sentences that contain either expected or surprising event boundaries have higher xReact and oReact, which are emotional responses either by the subject or by others to an event. For example, this is the case for the surprising and emotional event boundary "I remember it was like the 3rd or 4th game when something bad happened." This suggests that event boundaries are more likely when a sentence is more emotionally charged, echoing work by Dunsmoor et al. (2018) on how event segmentation is particularly frequent when the emotion of fear is triggered.

6 Comparison with Story Cloze Test

To better understand how surprising event boundaries relate to deviation from commonsense reasoning, we compare our Event Boundary Detection Task to the ROC Story Cloze Test (Mostafazadeh et al., 2016). This Test involves identifying whether a candidate ending sentence follows commonsense (*commonsense ending*) or deviates from commonsense (*nonsense ending*) given the first four sentences of a short story. Deviation from commonsense reasoning is one factor that can cause surprise (Sap et al., 2019a) and therefore comparing our task to the ROC Story Cloze Test can allow us to potentially isolate deviations from commonsense from other factors that can cause surprise. The ROC Story Cloze Test dataset contains 3142 samples with 1571 commonsense endings and 1571 nonsense endings.³ We train a separate Event Boundary Detector model on the ROC Story Cloze Test, using the same experimental setup as for event boundary detection, except the loss function; we use the cross-entropy loss since only one label is available for each sample.⁴

³We use the Winter 2018 version, which contains a dev and a test set. As in previous work (Schwartz et al., 2017), we train our model on the dev portion.

⁴Training takes 20 minutes on an Nvidia P100 GPU.

	overall F1	nonsense ending F1	commonsense ending F1
Event Detector w RoBERTa			
- ALLTOCURRENT	87.9	87.8	88.0
- PREVIOUSONLY	87.6	87.3	87.8
- SEQUENTIAL	87.3	87.1	87.5
RoBERTa	87.7	87.6	87.8

Table 3: ROC Story Cloze Test

Performance of Event Detector on ROC Story Cloze Test Compared to the Event Boundary Detection task, models perform significantly better on the ROC Story Cloze Test (highest F1 = 78.0% vs. 87.9%, $p < 0.001$ based on a two-tailed t-test, as observed in Tables 2 and 3). While the tasks are not directly comparable due to the inherent subjectivity of the Event Boundary Detection Task, the higher performance on the ROC Story Cloze Test suggests that identifying surprising, expected or no event boundaries may be more challenging than identifying commonsense or nonsense endings. Our commonsense and narrative features also do not seem to significantly improve upon RoBERTa’s performance in the ROC Story Cloze Test (+0.2% F1). This indicates that detecting whether a story ending follows commonsense can be effectively approached using RoBERTa alone, making it relatively easier to tackle.

7 Conclusion

We tackle the task of identifying event boundaries in stories. We propose a model that combines predictions made using commonsense and narrative features with a RoBERTa classifier. We found that integrating commonsense and narrative features can significantly improve the prediction of surprising event boundaries through detecting violations to commonsense relations (especially relating to the absence of causality), low likelihood events, and changes in emotional valence. Our model is capable in detecting event boundaries with high annotator agreement but limited in detecting those with lower agreement. Compared to identifying commonsense and nonsense story endings in Story Cloze Test, our task is found to be more challenging. Our results suggest that considering commonsense knowledge and narrative features can be a promising direction towards characterizing and detecting event boundaries in stories.

637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691

References

Wasi Uddin Ahmad, Nanyun Peng, and Kai-Wei Chang. 2021. [Gate: Graph attention transformer encoder for cross-lingual relation and event extraction.](#)

Christopher Baldassano, Uri Hasson, and Kenneth A. Norman. 2018. [Representation of real-world event schemas during narrative perception.](#) *The Journal of Neuroscience*, 38(45):9689–9699.

Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online. Association for Computational Linguistics.

Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Çelikyilmaz, and Yejin Choi. 2019. [Comet: Commonsense transformers for automatic knowledge graph construction.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Luis Emilio Bruni, Sarune Baceviciute, and Mohammed Arief. 2014. [Narrative cognition in interactive systems: Suspense-surprise and the p300 erp component.](#) In *Interactive Storytelling*, pages 164–175, Cham. Springer International Publishing.

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.

Yubo Chen, Shulin Liu, Xiang Zhang, Kang Liu, and Jun Zhao. 2017. [Automatically labeled data generation for large scale event extraction.](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 409–419, Vancouver, Canada. 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.

Joseph E. Dunsmoor, Marijn C. W. Kroes, Caroline M. Moscatelli, Michael D. Evans, Lila Davachi, and Elizabeth A. Phelps. 2018. [Event segmentation protects emotional memories from competing experiences encoded close in time.](#) *Nature Human Behaviour*, 2(4):291–299.

Meadhbh I. Foster and Mark T. Keane. 2015. [Predicting surprise judgments from explanation graphs.](#)

Saadia Gabriel, Chandra Bhagavatula, Vered Shwartz, Ronan Le Bras, Maxwell Forbes, and Yejin Choi. 2021. [Paragraph-level commonsense transformers with recurrent memory.](#) *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12857–12865.

Sarik Ghazarian, Zixi Liu, Akash S M, Ralph Weischedel, Aram Galstyan, and Nanyun Peng. 2021. [Plot-guided adversarial example construction for evaluating open-domain story generation.](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4334–4344, Online. Association for Computational Linguistics.

Jonathan Gordon and Benjamin Van Durme. 2013. [Reporting bias and knowledge acquisition.](#) In *Proceedings of the 2013 Workshop on Automated Knowledge Base Construction, AKBC '13*, page 25–30, New York, NY, USA. Association for Computing Machinery.

Arthur C Graesser, Scott P Robertson, and Patricia A Anderson. 1981. [Incorporating inferences in narrative representations: A study of how and why.](#) *Cognitive Psychology*, 13(1):1–26.

Anna Jafarpour, Elizabeth A Buffalo, Robert T Knight, and Anne GE Collins. 2019a. [Event segmentation reveals working memory forgetting rate.](#) *Available at SSRN 3614120*.

Anna Jafarpour, Sandon Griffin, Jack J Lin, and Robert T Knight. 2019b. [Medial orbitofrontal cortex, dorsolateral prefrontal cortex, and hippocampus differentially represent the event saliency.](#) *Journal of cognitive neuroscience*, 31(6):874–884.

CA Kurby and JM Zacks. 2009. [Segmentation in the perception and memory of events.](#) *Trends in cognitive sciences*.

Christopher A Kurby and Jeffrey M Zacks. 2008. [Segmentation in the perception and memory of events.](#) *Trends in cognitive sciences*, 12(2):72–79.

Qi Li, Heng Ji, and Liang Huang. 2013. [Joint event extraction via structured prediction with global features.](#) In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 73–82, Sofia, Bulgaria. 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.](#)

Lara Martin, Prithviraj Ammanabrolu, Xinyu Wang, William Hancock, Shruti Singh, Brent Harrison, and Mark Riedl. 2018. [Event representations for automated story generation with deep neural nets.](#) *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).

749	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 <i>Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 839–849, San Diego, California. Association for Computational Linguistics.	804	Corby Rosset. 2020. Turing-nlg: A 17-billion-parameter language model by microsoft .	805
752		806	Marie-Laure Ryan. 2010. Narratology and cognitive science: A problematic relation . <i>Style</i> , 44(4):469–495.	807
753		808		808
754		809	Maarten Sap, Eric Horvitz, Yejin Choi, Noah A Smith, and James W Pennebaker. 2020. Recollection versus imagination: Exploring human memory and cognition via neural language models . In <i>ACL</i> .	810
755		811		811
756		812		812
757		813	Maarten Sap, Anna Jafarpour, Yejin Choi, Noah A. Smith, James W. Pennebaker, and Eric Horvitz. 2022. Computational lens on cognition: Study of autobiographical versus imagined stories with large-scale language models .	814
758		815		815
759	Nasrin Mostafazadeh, Aditya Kalyanpur, Lori Moon, David Buchanan, Lauren Berkowitz, Or Biran, and Jennifer Chu-Carroll. 2020. GLUCOSE: Generalized and Contextualized story explanations . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 4569–4586, Online. Association for Computational Linguistics.	816		816
760		817		817
761		818	Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. 2019a. Atomic: An atlas of machine commonsense for if-then reasoning . <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , 33(01):3027–3035.	819
762		820		820
763		821		821
764		822		822
765		823		823
766		824		824
767	Aida Nematzadeh, Kaylee Burns, Erin Grant, Alison Gopnik, and Tom Griffiths. 2018. Evaluating theory of mind in question answering . In <i>Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing</i> , pages 2392–2400, Brussels, Belgium. Association for Computational Linguistics.	825	Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019b. Social IQa: Commonsense reasoning about social interactions . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.	826
768		827		827
769		828		828
770		829		829
771		830		830
772		831		831
773		832		832
774	Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences . <i>Transactions of the Association for Computational Linguistics</i> , 7:677–694.	833		833
775		834	R.C. Schank and R. Abelson. 1977. <i>Scripts, Plans, Goals, and Understanding</i> . Hillsdale, NJ: Earlbaum Assoc.	835
776		836		836
777		837	Roy Schwartz, Maarten Sap, Ioannis Konstas, Li Zilles, Yejin Choi, and Noah A Smith. 2017. The effect of different writing tasks on linguistic style: A case study of the roc story cloze task . In <i>CoNLL</i> .	838
778	Kyle A. Pettijohn and Gabriel A. Radvansky. 2016. Narrative event boundaries, reading times, and expectation . <i>Memory & Cognition</i> , 44(7):1064–1075.	839		839
779		840		840
780		841	Matthew Sims, Jong Ho Park, and David Bamman. 2019. Literary event detection . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 3623–3634, Florence, Italy. Association for Computational Linguistics.	842
781	Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.	843		843
782		844		844
783		845		845
784	Gabriel A. Radvansky, Andrea K. Tamplin, Joseph Armendarez, and Alexis N. Thompson. 2014. Different kinds of causality in event cognition . <i>Discourse Processes</i> , 51(7):601–618.	846	Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. Mpnet: Masked and permuted pre-training for language understanding . <i>arXiv preprint arXiv:2004.09297</i> .	847
785		848		848
786		849		849
787		850	David J. Townsend. 2018. Stage salience and situational likelihood in the formation of situation models during sentence comprehension . <i>Lingua</i> , 206:1–20.	851
788	Hannah Rashkin, Asli Celikyilmaz, Yejin Choi, and Jianfeng Gao. 2020. PlotMachines: Outline-conditioned generation with dynamic plot state tracking . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 4274–4295, Online. Association for Computational Linguistics.	852		852
789		853	Ercenur Ünal, Yue Ji, and Anna Papafragou. 2019. From event representation to linguistic meaning . <i>Topics in Cognitive Science</i> , 13(1):224–242.	854
790		855		855
791		856	Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus .	857
792		858		858
793				
794				
795	Hannah Rashkin, Maarten Sap, Emily Allaway, Noah A. Smith, and Yejin Choi. 2018. Event2mind: Commonsense inference on events, intents, and reactions . In <i>ACL</i> .			
796				
797				
798				
799	Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing</i> . Association for Computational Linguistics.			
800				
801				
802				
803				

- 859 Timothy D. Wilson and Daniel T. Gilbert. 2008. [Explaining away: A model of affective adaptation](#). *Perspectives on Psychological Science*, 3(5):370–386.
860
861 PMID: 26158955.
862
- 863 Lili Yao, Nanyun Peng, Weischedel Ralph, Kevin
864 Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-
865 and-write: Towards better automatic storytelling. In
866 *The Thirty-Third AAAI Conference on Artificial In-*
867 *telligence (AAAI-19)*.
- 868 Jeffrey M. Zacks. 2020. [Event perception and mem-](#)
869 [ory](#). *Annual Review of Psychology*, 71(1):165–191.
870 PMID: 31905113.
- 871 Jeffrey M Zacks, Nicole K Speer, Khena M Swallow,
872 Todd S Braver, and Jeremy R Reynolds. 2007. Event
873 perception: a mind-brain perspective. *Psychologi-*
874 *cal bulletin*, 133(2):273.

875 **A Appendix**

876 **A.1 Atomic relations training details**

877 We used the train/dev/test splits from the original
878 Atomic dataset. Negative samples are created by
879 matching a Atomic event node to a correspond-
880 ing tail event node from another sample based
881 on the relationship involved. Sepcifically, nega-
882 tive sampling was performed on groups ([’xWant’,
883 ’oWant’, ’xNeed’, ’xIntent’],[’xReact’, ’oReact’,
884 ’xAttr’],[’xEffect’, ’oEffect’]) given that the tail
885 event nodes in each group are more similar, creat-
886 ing more discriminating negative samples, as in-
887 spired by Sap et al. (2019b). One negative sample
888 is introduced every nine positive samples, since
889 there are nine labels. We used a learning rate of
890 1e-4, batch size of 64, 8 epochs and AdamW op-
891 timizer. Training took 18 hours on a Nvidia P100
892 GPU.

893 **A.2 Glucose relations training details**

894 Because the Glucose dataset was not split initially,
895 we randomly split the dataset into train/dev/test
896 splits based on a 80/10/10 ratio. For each sample
897 in Glucose, annotations share similar head event
898 nodes in Dim-1 to 5 and similar tail event nodes
899 in Dim-6 to 10. Therefore, our negative sampling
900 strategy for Dim-1 to 5 involves randomly choosing
901 a tail node from Dim-6 to 10 and vice-versa. As a
902 result, one negative sample is introduced every five
903 samples. During training, we used a learning rate
904 of 1e-4, batch size of 64, 8 epochs and AdamW
905 optimizer. Training took 15 hours on a Nvidia P100
906 GPU.

907 **A.3 Realis training details**

908 We used the train/dev/test split from the Realis
909 dataset. During training, we used the AdamW opti-
910 mizer, a learning rate of 2e-5, 3 epochs and batch
911 size of 4, as inspired by (Sap et al., 2020). Training
912 took 1 hour on a Nvidia P100 GPU.

913 **A.4 Sequentiality experimental details**

914 GPT2-small was accessed from HuggingFace
915 Transformers library and used without further fine-
916 tuning. It has 125M parameters, a context window
917 of 1024, hidden state dimension of 768, 12 heads
918 and dropout of 0.1.

919 **A.5 SimGen experimental details**

920 We used the Turing-NLG model without further
921 fine-tuning. The model has 17B and we used it

with top-p sampling (top-p=0.85), temperature=1.0
and max sequence length of 64 tokens. MPnet-
base model was accessed from the Sentence-BERT
library (Reimers and Gurevych, 2019) and used
without further fine-tuning.

927 **A.6 Event Boundary Detection Model** 928 **training details**

AdamW optimizer was used with $\alpha = 5 * 10^{-6}$, fol-
929 lowing a uniform search using F1 as the criterion at
930 intervals of $\{2.5, 5, 7.5, 10\} * 10^n; -6 \leq n \leq -3$.
931 Learning rate was linearly decayed (8 epochs) with
932 100 warm-up steps. Batch size of 16 was used. Val-
933 idation was done every 0.25 epochs during training.
934 Training each model took around 30 minutes on an
935 Nvidia P100 GPU.