

# Evaluation of Unsupervised Entity and Event Saliency Estimation

**Jiaying Lu**

Department of Computer Science  
Emory University  
Atlanta, GA 30322 USA  
jiaying.lu@emory.edu

**Jinho D. Choi**

Department of Computer Science  
Emory University  
Atlanta, GA 30322 USA  
jinho.choi@emory.edu

## Abstract

Saliency Estimation aims to predict term importance in documents. Due to few existing human-annotated datasets and the subjective notion of saliency, previous studies typically generate pseudo-ground truth for evaluation. However, our investigation reveals that the evaluation protocol proposed by prior work is difficult to replicate, thus leading to few follow-up studies existing. Moreover, the evaluation process is problematic: the entity linking tool used for entity matching is very noisy, while the ignorance of event argument for event evaluation leads to boosted performance. In this work, we propose a light yet practical entity and event saliency estimation evaluation protocol, which incorporates the more reliable syntactic dependency parser. Furthermore, we conduct a comprehensive analysis among popular entity and event definition standards, and present our own definition for the Saliency Estimation task to reduce noise during the pseudo-ground truth generation process. Furthermore, we construct dependency-based heterogeneous graphs to capture the interactions of entities and events. The empirical results show that both baseline methods and the novel GNN method utilizing the heterogeneous graph consistently outperform the previous SOTA model in all proposed metrics.

## 1 Introduction

Automatic extraction of prominent information from text has gained significant interest due to the huge volume of articles produced every day. Among various discourse elements, entity and event play a central role to form the skeleton of the story plot. Although entity and event salient estimation is critical for many downstream tasks, we notice that few studies follow-up Xiong et al. (2018) and Liu et al. (2018). The reasons behind that are the replication crisis of previous approaches and flawed evaluation protocol.

As the notion of saliency is subjective and varies from person to person, few human-annotated datasets exist. Hence, previous studies (Dunietz and Gillick 2014; Xiong et al. 2018; Liu et al. 2018) typically generate pseudo-ground truth, following the assumption that an entity or event is considered salient if a summary written by humans includes the reference of entity or event. This practical assumption allows

researchers to create large-scale pseudo annotations. While some prior work reported state-of-the-art performance on their proposed evaluation metrics, the strict dependency of their pre-processing pipelines hindered the reproducibility. Moreover, the noise introduced by entity linking system used in entity saliency task and the ignorance of event argument in event saliency task led to debatable scores.

In this paper, we investigate flawed evaluation protocols of previous work and present our own evaluation protocol. The proposed evaluation protocol first defines the mention form and extent of entity and event by combining various popular standards. We then employ a light yet accurate pseudo-ground truth annotation process\* that leverages more reliable dependency parser (He and Choi 2020). Same ranking metrics on top-k salient entities and events are incorporated as evaluation metrics. We argue that the proposed evaluation protocol is easier to replicate by other researchers, and the pseudo ground-truth contains less noise. Moreover, we conduct extensive experiments to re-examine existing approaches and recently proposed Graph Convolutional Neural Networks models that have demonstrated their power in many NLP tasks. Thus, we show that both baseline and novel heterogeneous R-GCN (Schlichtkrull et al. 2018) methods perform better than the previously proposed models.

## 2 Entity and Event (Pseudo) Annotation

Various tasks in entity and event processing have been introduced by previous studies (Ruppenhofer et al. 2010; Minard et al. 2015; Moro and Navigli 2015; Mitamura, Liu, and Hovy 2019; Ji et al. 2019). However, many of them have defined their own terms of entities and events (Bies et al. 2016), making it difficult to decide which standard to follow. Table 1 compares the entity/event annotation from the previous works to the one produced by this work. Our dataset is much larger than the others because it is pseudo-generated. Manual annotation of these kinds requires intensive human labor with a solid linguistics background, which makes the process resource-consuming. Such a large amount of pseudo-annotation in our dataset allows researchers from diverse disciplines to perform big data analysis on entities and events occurring over a long period of time.

\*The pseudo annotation is available at <https://github.com/lujiaying/pseudo-ee-saliency-estimation>

| Dataset     | Typology                                    | Form    | Entities |
|-------------|---|---------|----------|
| ACE         | P/O/L/G/F, Vehicle, Weapon + 30 subtypes    | N/NE/NP | 56,898   |
| ECB+        | P/O/L/G/F, Artifact, Generic, Time, Vehicle | N/NE    | 17,187   |
| Rich ERE    | P/O/L/G/F + 20 subtypes                     | N/NE/NP | 5,873    |
| RED         | All entity types                            | N/NE    | 10,320   |
| <b>Ours</b> | All entity types                            | N/NE/NP | 144.3M   |

(a) Comparisons of entity annotation. Entities: the total number of annotated entities, P/O/L/G/F: Person/Organization/Location/Geopolitical entity/Facility, N/NE/NP: Noun/Named Entity/Noun Phrase.

| Dataset     | Typology $\Rightarrow$ Source  | Form       | Events   |
|-------------|--|------------|----------|
| TimeML      | Occurrence, Action, State, Aspectual, Perception, Reporting<br>$\Rightarrow$ 300 docs from news (Pustejovsky et al. 2003)  | V/N/J/C/PP | 7,571    |
| ACE         | Business, Conflict, Contact, Justice, Life, Movement, Personnel, Transaction<br>+ 33 subtypes $\Rightarrow$ 599 docs from news and broadcasts (Doddington et al. 2004)                                     | V/J        | 5,312 ✓  |
| ECB+        | Occurrence, State, Aspectual, Causation, Generic, Perception, Reporting<br>$\Rightarrow$ 982 docs from news (Cybulska and Vossen 2014)   | V/N/J      | 14,884 ✓ |
| Rich ERE    | Business, Conflict, Contact, Generic, Irrealis, Justice, Life, Manufacture,<br>Movement, Personnel, Transaction + 38 subtypes<br>$\Rightarrow$ 505 docs from news and discussion forums (Song et al. 2015) | V/N/J/R    | 2,933 ✓  |
| RED         | Occurrence, Action, State, Process<br>$\Rightarrow$ 95 docs from news and discussion forums (O’Gorman, Wright-Bettner, and Palmer 2016)  | V/N/J      | 8,731    |
| <b>Ours</b> | Occurrence, Action, State, Process<br>$\Rightarrow$ 1.65M docs from news and scientific articles   | V/N        | 58.9M ✓  |

(b) Comparisons of event annotation. Events: # of annotated events, ✓: events are annotated with their arguments, V/N/J/R/C/PP: predicative Verb/Noun/adjective/adverb/Clause/Preposition Phrase.

Table 1: Comparisons of entity and event annotation in the previous and our works.

The trend is clear that the more recent work it is, the more entity and event types are covered by the work. Our annotation is the most similar to the RED dataset (O’Gorman, Wright-Bettner, and Palmer 2016) that covers the broadest types of entities and events, although entities in RED are annotated on the heads unless they are named entities whereas they are annotated on the base Noun Phrases in our dataset, which is more explicit (Section 2). Moreover, events are annotated with their core arguments in our dataset, which are not present in RED. It is worth mentioning that RED comprises extra layers of annotation (e.g., polarity, modality) as well as predicative adjectives for events not included in our dataset, which we will explore in the future.

### Entity Definition

Inspired by definitions in previous work, we consider all **base Noun Phrases** (base NPs) as entity candidates, excluding eventive nouns. Although some base NPs, such as *president* or *two weeks*, may refer to multiple real-world objects which violate the rigid definition, this simplification is beneficial to pseudo annotation. Another advantage of using base NPs instead of higher-level NPs is the more fine-grained entity annotation.

### Event Definition

An event describes “*who did what to whom when and where*”. Therefore, the event trigger (*what*) itself can not represent a complete event. In our definition, event triggers and core arguments are essential components of events. **Verbs, Eventive Nouns** (deverbal nouns, proper names referring to historically significant events), phrase constituted by **Light Verb +**

**Noun** and predicative **Adjectives** are generally considered as **event triggers** in prior studies. For annotation simplicity, adjectives are not included in our definition. Since not all verbs and nouns are valid event triggers, we create a pre-defined vocabulary using following procedures:

1. We collect the candidate list of verbs and deverbal nouns from FrameNet (Baker, Fillmore, and Lowe 1998) and NomBank (Meyers et al. 2004) utilizing the same 569 frames generated by (Liu et al. 2018).
2. We then manually add valid head words (around 47) of proper names such as *epidemic*, *earthquake*, etc.
3. Similar to previous work, we remove auxiliary and copular verbs<sup>†</sup>, light verbs<sup>‡</sup>, and report verbs<sup>§</sup>, as they are rarely representative events.

This gives us a total of 2645 verb lemmas and 516 eventive nouns. Regarding core arguments, we consider entity and sub-event participants specifying *who* or *whom* is involved in the event.

### Pseudo Annotation

For articles in the Corpus, we annotate pseudo salient entities and events by extracting entity, event trigger, event core argument mentions in the human written abstracts according to the aforementioned definitions. Dependency parser is the

<sup>†</sup> Auxiliary and copular verbs include *appear*, *become*, *do*, *have*, *seem*, *be*.

<sup>‡</sup> Light verbs include *do*, *get*, *give*, *go*, *have*, *keep*, *make*, *put*, *set*, *take*.

<sup>§</sup> Report verbs include *argue*, *claim*, *say*, *suggest*, *tell*.

main tool we leverage to conduct the pseudo annotation. We believe the dependency parser is more reliable than the entity linker used in previous studies, since the SOTA model achieved above 97% Unlabeled Attachment Score (UAS) on Penn Treebank (He and Choi 2020) and the parser is easy to integrate with other NLP tools. The pseudo annotation process is as follows: First, sentences in the abstract are parsed by a syntactic dependency parser. Second, heuristics<sup>1</sup> are applied to identify *base NPs* and *verbs* upon POS taggers and dependency relations (Choi 2017). Then, all base NPs excluding ones that have eventive noun heads are extracted as entity candidates, while complement base NPs are event candidates. Finally, verbs or the action carrier verbs within verb phrases (i.e. discarding auxiliary verbs and/or raising verbs in verb phrases) in the pre-defined vocabulary are extracted as event trigger candidates. In the above eventive noun and verb cases, only the head part would be annotated:

*Allen was walking quickly to the mall.*  
*This terrible war could have ended in a month.*

Besides, in the *light verbs + noun* structure, only the nominal part (typically eventive noun) would be annotated:

*John made a call to Mary.*  
*John gave an interview to Mary.*

For event core arguments, we consider entities or sub-events that have *nominal subject*, *object* or *dative* relations with event trigger. For instance, *John* and *Mary* are annotated as core arguments in above cases.

### 3 Proposed Approach

The salience identifying ability of humankind depends significantly on relational structures of information. Intuitively, we construct an undirected graph per input article which distills the important knowledge for entities and events salience estimation. The graph  $\mathcal{G}$  is a heterogeneous graph that consists of two types of nodes: entity nodes, and event nodes. The node types can be further expanded by including coreferential cluster nodes, sentence nodes, and others when necessary. The details about node candidate generation is described in subsection Candidate Generation Module. Then the graph construction module collects these generated entity and event nodes and connects them accordingly. Different from most previous work which considers the graph as a fully-connected graph, edges are added according to the dependency tree arcs which lead the plot graph  $\mathcal{G}$  a partially-connected graph. Multiple types of edges including dependency edges, coreferential edges, and inter-sentence edges are added to better reflect semantic and syntactic relations between nodes.

Once the nodes and edges are established, edge weights are critical for salience node prediction. For initial edge weights, we design multiple mechanisms to utilize the frequency, position, and other information as much as possible. The weighted graph can be directly used for classic graph ranking algorithms. However, heuristic-based weights design can be tricky and hard to generalize. To overcome the disadvantages

<sup>1</sup>heuristics for base NP using context-free grammar:  $NP \rightarrow DT \bar{N}$ ;  $\bar{N} \rightarrow NN$ ;  $\bar{N} \rightarrow NN \bar{N}$ ;  $\bar{N} \rightarrow JJ \bar{N}$ ;  $\bar{N} \rightarrow \bar{N} \bar{N}$

above, we train a heterogeneous R-GCN (Schlichtkrull et al. 2018) model to learn the node salience by feeding top frequent entities and events as the alternative training data.

### Candidate Generation Module

Figure 1 shows an example of entity and event candidate spans generated by a dependency tree. For simplicity, each span in figure 1 is represented by the dependent head. The extraction result can be regarded as the remaining dependency tree after removing auxiliary words.

As mentioned in Section 2, entities and events have their own syntactic properties:

- 1) Entities are generally *noun phrases*.
- 2) Events are *verbal predicates* or *nominal predicates*.

It is then intuitively to derive context-free grammar rules which leverage the POS tagging and dependency relation information to extract entities and events. Since the extraction process is relatively deterministic, the result is reliable. Our implementation uses a BERT-based Dependency Parser from He and Choi (2020) and can be extended using any dependency parser.

[Kobe Bryant] was killed [Sunday] when [the helicopter] [he] was traveling in crashed and burst into [flames] in the [hills] above [Calabasas].

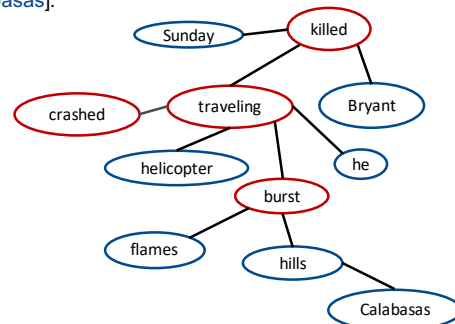


Figure 1: Dependency-based Candidate Generation

### Graph Construction Module

After all entity and event candidates are extracted, it is natural to construct a graph to capture the relation structure of information. In addition to the intra-sentence dependency relations from dependency parsing, inter-sentence level relations such as adjacent sentence syntactic roots, coreference resolution is also introduced in our graph construction model. The graph ranking algorithm is then deployed for salience identifying. In the following subsections, we will describe the proposed construction algorithm for the graph.

**Dependency-based Graph Construction** Figure 2 shows a dependency-based graph for two sentences about Kobe Bryant’s helicopter crash. Adjacent noun phrases edges, coreferential edges, and adjacent sentence syntactic roots edges are further included in the graph to capture more information. First, the initial dependency-based graph is already established after the candidate extraction phase. The edge  $w_{ij}^{dep}$  between two nodes  $v_i$  and  $v_j$  are weighted according




Figure 2: Dependency-based Graph with extra edges introduced. Dependency, Coreferential, and Adjacent sentence edges are denoted by blue, purple and red edges, respectively.

to the sum of the inverse tree distance between each span headword of  $hw_i$  and  $hw_j$ , as equation 1 denoted. The tree distance is defined as the minimum number of edges to be traversed to reach one node from another. It is also worth noting that in dependency-based graph nodes that share the same span have not been merged yet.

$$w_{ij}^{dep} = \frac{1}{tree\_dist(hw_i, hw_j)}, \forall e_{ij} \in E_{dep} \quad (1)$$

Next, adjacent NPs and syntactic roots are connected accordingly. Edges between Adjacent NPs ( $E_{NPs}$ ) reflect the information propagation between neighbor entities and contain both intra- and inter-sentence level information. The edge weights  $w^{NPs}$  are set by the inverse of tree distances as well. On the other hand, edges between adjacent syntactic roots  $E_{roots}$  reflect the information propagation between neighbor sentences. Since the dependency tree only provides within sentence distances, we can assign a constant to  $w^{roots}$  (in practice, we choose 1) as the virtual distance for each edge of  $E_{roots}$ . Therefore, equation 1 can be expanded to include edges in  $E_{NPs}$  and  $E_{roots}$ .

Moreover, spans in same coreferential clusters are connected. The coreferential clusters can be obtained from either ground truth of the corpus or system-generated result. The edge weights  $w^{coref}$  of  $E_{coref}$  can also be assigned to a constant (we still use 1 in experiments) because spans in one cluster are equally important. Finally, NP spans consisting of the same tokens within one article  $\mathcal{D}$  would then be merged into one node  $x$  in the plot graph  $\mathcal{G}$ , except single pronoun spans. Therefore,

$$x_i = \{s_{i,j}, \dots, s_{i,k}\} \text{ where } s_{i,j} = s_{i,k} \quad (2)$$

In equation 2,  $s_{i,j}$  denotes the  $j$ th span of node  $x_i$ . The overall weights between node  $x_i$  and  $x_j$  then become the sum of different types of weights, which is shown in Equation 3.  $w_{kl}^{dep}$ ,  $w_{kl}^{NPs}$ ,  $w_{kl}^{roots}$ ,  $w_{kl}^{coref}$  represent edge weights from dependency tree, adjacent Nps, adjacent syntactic roots and conference resolution clusters between  $span_{i,k}$  and  $span_{j,l}$ , respectively.

$$w_{ij} = \sum_{s_{i,k} \in x_i} \sum_{s_{j,l} \in x_j} \frac{1}{w_{kl}^{dep} + w_{kl}^{NPs} + w_{kl}^{roots} + w_{kl}^{coref}} \quad (3)$$

After the edge weights are set, classic unsupervised graph ranking such as TextRank (Mihalcea and Tarau 2004) algorithms can be directly employed. The importance of nodes in the graph is calculated and then used as indicators of whether entities or nodes are salient.

## Graph Neural Network upon Heterogeneous Graph

The advantage of classic graph ranking methods is that no training data is required. However, the design of edge weights is tricky and may significantly impact the final performance. A learning-based model is preferable to let the model automatically learn the importance of different edge types given source and target vertices. Therefore, a relational GCN (Schlichtkrull et al. 2018) is adapted. The key idea of relation GCN is that edge weights are determined by edge types and can be learned through a training process. Equation 4 illustrates the vertex representation update in R-GCN, where  $c_{i,r}$  denotes a normalization constant and  $W_r$  denotes the learnable weight for edge type  $r$ . The choice of initial node embedding is flexible. Bag-of-words embedding, static word embedding, or contextual word embedding can be used, with POS tagging, named entity results as optional extra features. According to the graph construction, edge types include dependency relations, adjacent noun phrase, adjacent sentence root, and coreferent vertices. Since the constructed graph is a heterogeneous graph, the edge types are further categorized by the source and target vertex types.

$$h_i^{l+1} = \sigma(W_0^l h_i^l + \sum_r \sum_j \frac{1}{c_{i,r}} W_r^l h_j^l) \quad (4)$$

Although no ground-truth entities and events are provided in datasets, we choose top-frequent entities and events as alternative training samples. Thus, the R-GCN model is trained in a supervised node classification setting. During inference, heterogeneous graphs are constructed for input text using the same methods and then used to predict salient nodes.

## 4 Experiments

### Datasets

| Dataset | # of Docs in Trn /Dev /Tst | # of E&E | avg of E&E |
|---------|----------------------------|----------|------------|
| NYT     | 526K /64K /64K             | 98M /39M | 150 /60    |
| SS      | 800K /100K /100K           | 46M /19M | 46 /20     |

Table 2: Statistics for New York Times and Semantic Scholar datasets. (Trn, Dev, Tst: Train, Development, Test set. E&E: Entity and Event.)

We evaluate the proposed protocol on a news corpus and a scientific publication corpus as follows.

- **New York Times** (Dunietz and Gillick 2014) provides 654K online news articles that contains experts written summarizes. The entities and events annotated on the summarizes are considered as salient.
- **Semantic Scholar** (Xiong, Power, and Callan 2017) contains one million randomly sampled scientific publications. Since only the abstracts and titles of papers are available, we consider entities that occurred on titles as salient.

| Method    | New York Times |              |              |              | Semantic Scholar |              |              |              |
|-----------|----------------|--------------|--------------|--------------|------------------|--------------|--------------|--------------|
|           | F1@1           | F1@3         | F1@5         | F1@10        | F1@1             | F1@3         | F1@5         | F1@10        |
| TextRank  | 3.69           | 8.22         | 10.60        | 13.17        | 5.46             | 8.60         | 9.43         | 9.49         |
| KCE       | 5.46           | 11.64        | 14.52        | 16.57        | 4.83             | 7.82         | 8.41         | 7.77         |
| Frequency | <b>7.02</b>    | 14.00        | 17.69        | 22.40        | <b>10.78</b>     | <b>16.49</b> | <b>17.11</b> | <b>15.08</b> |
| Location  | 3.94           | 10.26        | 16.40        | <b>23.86</b> | 4.65             | 12.08        | 13.76        | 13.45        |
| R-GCN     | 6.65           | <b>14.03</b> | <b>18.10</b> | 22.72        | 9.66             | 15.39        | 16.35        | 14.81        |

Table 3: Entity salience estimation F1 scores (in percentage).

## Proposed Evaluation Metrics

In the paper which Dunietz and Gillick (2014) first introduced precision, recall and  $F1$  as evaluation metrics since the task was formulated as a binary classification problem. In the follow-up studies, Xiong et al. (2018) and Liu et al. (2018) argued that the importance of entities and events was on a continuous scale. They then proposed evaluating the salience task by ranking metrics: precision, recall at top 1, 5 and 10 salient entities and events. Although the ranking metrics can reveal the ability of salience identification, the way they calculated metrics was doubtful. For instance, if one mention  $x_i = \{t_a, t_b, \dots, t_z\}$  occurred in the expert-written abstract was accidentally linked to an entity  $E_a$  and the mention with exactly same tokens  $x_j = \{t_a, t_b, \dots, t_z\}$  in expert-written abstract was linked to another entity  $E_b$ ,  $x_i$  would be considered as non-salient. This violated the assumption that "an entity is salient if a reference summary includes it."

**(Macro-averaged) P, R, F1 @Top-k** To address the disadvantages and the inconsistency of previous evaluations, we propose different metrics for entities and events to better measure the salience estimation ability. Both metrics neglect punctuation marks and articles (*a, an, the*). These metrics highly rely on the pseudo annotated salient entities and events. For entity salience estimation, duplicated mentions of pseudo annotated salient entities and candidate salient entities are first merged separately. Then, precision (P), recall (R), F1 score are calculated between pseudo ground-truth and prediction using *exact match*. For event salience estimation, P, R, F1 are conducted between pseudo annotated event triggers and predicted triggers; meanwhile, between pseudo annotated event triggers with corresponding event arguments and predicted triggers with corresponding arguments.

In conclusion, our proposed evaluation protocol is superior to previous protocols for following reasons.

1) the flexibility and reproducibility for using any NLP tools to extract entity and event mentions.

2) the finer definition of a complete event that not only considers event trigger but also event arguments.

3) the less noisy dependency parser for pseudo ground-truth annotation.

## Compared Methods

**Previous Methods** In the experiments, the classic unsupervised method TextRank (Mihalcea and Tarau 2004) and the previous SOTA supervised method Kernel-based Centrality Estimation (KCE) (Liu et al. 2018) are analyzed, along with multiple baseline models and our novel unsupervised GNN based model. It is worth noting that KCE only considered

event triggers while producing salient events, and KCE did not conduct event salience estimation on Semantic Scholar. Therefore, numbers for corresponding columns in the performance table are blank.

**Baseline Methods** For both Frequency and Location methods, we first generate candidate entities and events by employing the same pipelines of heterogeneous graph construction. The frequency method then estimates the salience of an entity by its lemma frequency, and the salience of an event by its frequency of full lemma consisting of trigger and arguments. If two candidates have same frequencies, the one that occurred earlier would be ranked higher. On the other hand, the Location method essentially estimates the term salience by its occurrence position in the text, the earlier the higher.

**Proposed Heterogeneous GCN model** We construct a graph per article with 2 node types (entity, event), 92 edge types (bidirectional; including self edge, adjacent node edge, adjacent sentence edge, coreference edge, syntactic dependency edge, same headword edge). Current initial node embedding of is Glove embedding (2014) concatenating random initialized pos-tagging embedding, and then adding sinusoidal positional embedding. A heterogeneous R-GCN that learns different edge weights for each unique (source node, edge, target node) is applied to capture the characteristics of each node/edge type and to model the interaction between nodes. We treat the top-10 first occurred entities and events as salient terms. Therefore, the heterogeneous R-GCN is trained in the unsupervised setting since no ground-truth data is provided.

## Result Analysis

Table 3 shows the entity salience estimation performance, while Table 4 shows the event salience estimation performance on trigger only and trigger with arguments. On **entity salience estimation**, we observe that the Frequency baseline method and our novel heterogeneous R-GCN approach outperforms classic unsupervised TextRank and SOTA supervised KCE significantly for all  $F1@1,3,5,10$  metrics in both New York Times and Semantic Scholar datasets. On the other hand, the Location baseline method has better performance for all metrics except  $F1@1$ . Although the proposed R-GCN approach does not achieve the best performance in all metrics, the improvement over SOTA methods is consistent (1.19% to 6.55% in NYT, 4.83% to 7.94% in SS). It is also worth noting that R-GCN has not achieved the most competitive performance for every metric in Semantic Scholar dataset because titles are considered as the evaluation reference in this setting and titles are generally short sentences containing words occurring repeatedly in documents. On

| Method    | New York Times |              |              |              |              |              |              |              | Semantic Scholar |              |              |              |              |              |              |              |
|-----------|----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|           | F1@1           |              | F1@3         |              | F1@5         |              | F1@10        |              | F1@1             |              | F1@3         |              | F1@5         |              | F1@10        |              |
|           | Trig           | +Arg         | Trig         | +Arg         | Trig         | +Arg         | Trig         | +Arg         | Trig             | +Arg         | Trig         | +Arg         | Trig         | +Arg         | Trig         | +Arg         |
| TextRank  | 13.72          | 10.82        | 21.16        | 16.01        | 22.84        | 16.96        | 22.85        | 16.44        | 19.34            | 17.52        | 19.49        | 16.48        | 17.12        | 13.83        | 13.72        | 10.70        |
| KCE       | 13.11          | N/A          | 15.42        | N/A          | 15.65        | N/A          | 15.65        | N/A          | N/A              | N/A          | N/A          | N/A          | N/A          | N/A          | N/A          | N/A          |
| Frequency | <b>17.19</b>   | <b>12.84</b> | 24.85        | 18.81        | 26.19        | 19.82        | 24.88        | 18.71        | <b>23.92</b>     | <b>21.53</b> | <b>23.08</b> | <b>19.73</b> | <b>19.29</b> | 16.06        | 14.53        | 11.75        |
| Location  | 15.36          | 9.82         | <b>26.79</b> | 18.35        | <b>30.58</b> | <b>21.44</b> | 30.45        | <b>21.70</b> | 16.61            | 15.02        | 19.33        | 16.43        | 17.91        | 14.73        | 14.62        | 11.66        |
| R-GCN     | 16.15          | 12.08        | 25.59        | <b>18.98</b> | 28.12        | 20.91        | <b>28.47</b> | 20.92        | 20.79            | 19.28        | 21.57        | 18.88        | 19.19        | <b>16.15</b> | <b>15.00</b> | <b>12.06</b> |

Table 4: Event salience estimation F1 scores (in percentage). *Trig*, *+Arg* denotes evaluation on trigger only and trigger with corresponding arguments, respectively. (Event detection was not conducted by KCE on Semantic Scholar as authors stated.)

**event salience estimation**, we observe similar performance margins for approaches upon the constructed graph, indicating the effectiveness in leveraging the distilled structured relations as information to identify salient entities and events. Among the approaches utilizing the constructed graphs, baseline methods have advantages in detecting small numbers of top salient terms, while the R-GCN model is an expert at capturing bigger numbers of top salient terms.

## 5 Conclusion

In this paper, we revisit the evaluation metrics of the salience task, and reveal the drawbacks of previous work. To address the experiment reproducibility, noise from entity linker, and duplicated event trigger issues, we propose a light yet practical evaluation protocol that relies on the stable syntactic dependency parser. The empirical experiment result shows that our approaches are better than the SOTA approaches in the proposed metrics, which indicates the effectiveness of the constructed heterogeneous graph. Furthermore, we provide quantitative results on the complete event rather than the mere event trigger. In future work, we will improve our approach by employing an EM-based unsupervised algorithm, aiming to surpass the baseline approaches.

## References

- Baker, C. F.; Fillmore, C. J.; and Lowe, J. B. 1998. The berkeley framenet project. In *COLING'1998*.
- Bies, A.; Song, Z.; Getman, J.; Ellis, J.; Mott, J.; Strassel, S.; Palmer, M.; Mitamura, T.; Freedman, M.; Ji, H.; and O’Gorman, T. 2016. A comparison of event representations in DEFT. In *Proceedings of the Fourth Workshop on Events*.
- Choi, J. 2017. Deep dependency graph conversion in english. In *Proceedings of the 15th International Workshop on Treebanks and Linguistic Theories*.
- Cybulska, A., and Vossen, P. 2014. Using a sledgehammer to crack a nut? lexical diversity and event coreference resolution. In *LREC'14*.
- Doddington, G. R.; Mitchell, A.; Przybocki, M. A.; Ramshaw, L. A.; Strassel, S. M.; and Weischedel, R. M. 2004. The automatic content extraction (ace) program-tasks, data, and evaluation. In *LREC*.
- Dunietz, J., and Gillick, D. 2014. A new entity salience task with millions of training examples. In *EACL'14*.
- He, H., and Choi, J. D. 2020. Establishing Strong Baselines for the New Decade: Sequence Tagging, Syntactic and Semantic Parsing with BERT. In *FLAIRS'20*.
- Ji, H.; Sil, A.; Dang, H. T.; Soboroff, I.; Nothman, J.; and Hub, S. I. 2019. Overview of TAC-KBP 2019 Fine-grained Entity Extraction. In *TAC:KBP'19*.
- Liu, Z.; Xiong, C.; Mitamura, T.; and Hovy, E. 2018. Automatic event salience identification. In *EMNLP'18*.
- Meyers, A.; Reeves, R.; Macleod, C.; Szekely, R.; Zielinska, V.; Young, B.; and Grishman, R. 2004. The NomBank project: An interim report. In *HLT-NAACL 2004*.
- Mihalcea, R., and Tarau, P. 2004. TextRank: Bringing order into text. In *EMNLP'04*.
- Minard, A.-L.; Speranza, M.; Agirre, E.; Aldabe, I.; van Erp, M.; Magnini, B.; Rigau, G.; and Urizar, R. 2015. TimeLine: Cross-document event ordering. In *SemEval 2015*.
- Mitamura, T.; Liu, Z.; and Hovy, E. 2019. Events Detection, Coreference and Sequencing: What’s next? In *TAC:KBP'17*.
- Moro, A., and Navigli, R. 2015. SemEval-2015 task 13: Multilingual all-words sense disambiguation and entity linking. In *SemEval 2015*.
- O’Gorman, T.; Wright-Bettner, K.; and Palmer, M. 2016. Richer event description: Integrating event coreference with temporal, causal and bridging annotation. In *Proceedings of the 2nd Workshop on Computing News Storylines*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *EMNLP'14*.
- Pustejovsky, J.; Castano, J. M.; Ingria, R.; Sauri, R.; Gaizauskas, R. J.; Setzer, A.; Katz, G.; and Radev, D. R. 2003. Timeml: Robust specification of event and temporal expressions in text. *New directions in question answering*.
- Ruppenhofer, J.; Sporleder, C.; Morante, R.; Baker, C.; and Palmer, M. 2010. SemEval-2010 task 10: Linking events and their participants in discourse. In *Proceedings of the 5th International Workshop on Semantic Evaluation*.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Van Den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling relational data with graph convolutional networks. In *ESWC*.
- Song, Z.; Bies, A.; Strassel, S.; Riese, T.; Mott, J.; Ellis, J.; Wright, J.; Kulick, S.; Ryant, N.; and Ma, X. 2015. From light to rich ERE: Annotation of entities, relations, and events. In *Proceedings of the The 3rd Workshop on EVENTS*.
- Xiong, C.; Liu, Z.; Callan, J.; and Liu, T.-Y. 2018. Towards better text understanding and retrieval through kernel entity salience modeling. In *ACM SIGIR'18*.
- Xiong, C.; Power, R.; and Callan, J. 2017. Explicit semantic ranking for academic search via knowledge graph embedding. In *WWW'17*.