

Leveraging Contextual Information for Effective Entity Salience Detection

Rajarshi Bhowmik Marco Ponza Atharva Tendle Anant Gupta
Rebecca Jiang[†] Xingyu Lu Qian Zhao Daniel Preotiuc-Pietro

Bloomberg

{rbhowmik6, dpreotiucpie}@bloomberg.net

Abstract

In text documents such as news articles, the content and key events usually revolve around a subset of all the entities mentioned in a document. These entities, often deemed as salient entities, provide useful cues of the aboutness of a document to a reader. Identifying the salience of entities was found helpful in several downstream applications such as search, ranking, and entity-centric summarization, among others. Prior work on salient entity detection mainly focused on machine learning models that require heavy feature engineering. We show that fine-tuning medium-sized language models with a cross-encoder style architecture yields substantial performance gains over feature engineering approaches. To this end, we conduct a comprehensive benchmarking of four publicly available datasets using models representative of the medium-sized pre-trained language model family. Additionally, we show that zero-shot prompting of instruction-tuned language models yields inferior results, indicating the task's uniqueness and complexity.

1 Introduction

Many NLP studies have highlighted the importance of entities to understanding the semantics of a document (Wu et al., 2020b; Meij et al., 2012). Automatically identifying entities in unstructured text documents and linking them to an underlying knowledge base, such as Wikipedia, is one of the core NLP tasks, with multiple shared tasks (Tjong Kim Sang and De Meulder, 2003; Strauss et al., 2016), benchmarks (Hoffart et al., 2011; Hovy et al., 2006; Pradhan et al., 2013; Rijhwani and Preotiuc-Pietro, 2020; Derczynski et al., 2016), and studies (Kolitsas et al., 2018; Nguyen et al., 2014) dedicated to solving them.

Although an entity may play a crucial semantic role in document understanding, not all entities in

[†]Work was done while the author was affiliated with Bloomberg




Figure 1: An example of a document with salient and non-salient entities. Entity mentions are highlighted in text.

a text document play equal roles. Some entities are the central subjects or actors within a document, around which the content and the key events revolve. Others are mentioned only to provide additional context to the main event. For example, some entities may be actors in peripheral events, while others are deemed uninformative to the understanding of the document. Thus, *entity salience* in a text is defined as a binary or ordinal rating to quantify the extent to which a target entity is central to a given piece of text (Gamon et al., 2013; Dunietz and Gillick, 2014). Figure 1 provides an example text along with the mentioned entities and their salience. We note that the salience of an entity to a text is independent of the user's interest when reading or searching the document (Gamon et al., 2013), which is usually referred to as *entity relevance*. It is also distinct from *entity importance*, which quantifies the overall importance of the entity independent of the document. Automatically inferring entity salience was shown to aid search (Gamon et al., 2013), improve ranking results (Xiong et al., 2018), entity detection (Trani et al., 2018), and enable entity-centric applications such as entity-centric summarization (Maddela et al., 2022).

In this paper, we study the effectiveness of Transformer-based Pre-trained Language Models (PLMs) in the task of entity salience detection. Prior work on determining entity salience relied on heavy feature engineering to craft features explicitly covering relevant aspects, such as entity frequency (Dunietz and Gillick, 2014; Dojchinovski et al., 2016), position of entity mentions within a document (Dunietz and Gillick, 2014; Trani et al., 2018), relations to other entities (Trani et al., 2018), document features, such as its length (Gamon et al., 2013) and lexical features, such as the name of the entity or its context. Only a single recent work attempted to use PLMs in a pipeline which included key entity detection, albeit the scope of the evaluation was limited to a single high performing dataset (Zhao et al., 2021). In contrast, our proposed method uses a cross-encoder architecture where a target entity’s name or alias and its contextual mentions in a text document are encoded by a PLM encoder. The classifier uses the contextual representation and, optionally, positional information about the entity encoded through the decile position embedding vector of mentions to determine the salience score of a target entity.

We conduct experiments on four publicly available datasets, two of which were human annotated and two that were curated semi-automatically. We fine-tune several cross-encoders using PLMs and demonstrate that these yield consistent and significant improvements over feature-based methods, as well as prompting instruction-tuned PLMs. The latter shows the novelty and complexity of the task of entity salience detection, which requires the model to learn significant task-specific semantic knowledge for this natural language understanding task.

Our contributions in this paper are the following:

- We propose a cross-encoder style architecture with explicit encoding of position information for entity salience detection that shows consistent improvements of 7 – 24.4 F1 scores over previous feature engineering approaches.
- We establish a uniform benchmark of two human annotated and two semi-automatically curated datasets for the task of entity salience detection that we expect to be beneficial to future study of this task;
- A faceted analysis of the models’ predictive behaviour.

2 Related Work

Understanding the aboutness of a document is one of the long-standing goals of research in both Information Retrieval and Natural Language Processing (Gamon et al., 2013). Several types of approaches have been proposed, including extracting key-terms (Hulth, 2003; Mihalcea and Tarau, 2004), identifying latent topics (Blei et al., 2003), or generating text summaries (Erkan and Radev, 2004). There has been a recent focus in using entities to understand the content of a document. Towards this goal, the task of entity salience has been first described for web pages in (Gamon et al., 2013) and for news content in (Dunietz and Gillick, 2014). This task can be viewed as a restricted form of keyword or keyphrase extraction (Alami Merrouni et al., 2020) if salience is binary. For the rest of this study, we will use the concept of salience as described in (Gamon et al., 2013).

The salience labels for entities were obtained either by crowdsourcing labels from multiple raters to identify salient entities (Gamon et al., 2013; Dojchinovski et al., 2016; Trani et al., 2018; Madela et al., 2022) or by using proxies. For example, (Dunietz and Gillick, 2014) hypothesize that salient entities are those that appear in the article’s abstract. (Wu et al., 2020a) identifies an entity as salient if the Wikinews category that corresponds to the entity is also labeled as the category of the article.

Past studies mostly proposed machine learning methods to infer the salience of a given entity that relied on hand-crafted features. Features that can be computed from the target entity mentions and document alone can be categorized into the following: positional (e.g., position in the document, if entity is in the abstract) (Dunietz and Gillick, 2014), count-based (e.g., number of references to the entity) (Dunietz and Gillick, 2014; Wu et al., 2020a), local context (Trani et al., 2018), or global context (Ponza et al., 2019). Further, joint entity salience resolution can be performed by creating features using the entity graph (e.g., centrality in the entity graph) (Dunietz and Gillick, 2014; Trani et al., 2018). Finally, past work also showed that incorporating external knowledge about entities from knowledge bases can boost predictive performance (Dojchinovski et al., 2016).

Automatically inferring salience for entities can directly benefit multiple downstream applications, such as improving ranking results for queries containing entities (Xiong et al., 2018) or improv-

ing the performance of entity detection by joint modelling (Trani et al., 2018). Moreover, by inferring salience, new entity-centric applications can be built, such as highlighting salient entities in search (Gamon et al., 2013), improving the interpretability of news trends through salient entities (Ponza et al., 2021), or identifying entities for creating entity-centric summaries of news stories (Maddela et al., 2022; Hofmann-Coyle et al., 2022).

3 Problem Definition

We use the concept of salience as introduced in (Gamon et al., 2013): salient entities are entities explicitly mentioned in the document that are objectively important as a function of the structure of the text.

The goal of the salience model is to produce a single salience score $\psi(e)$ for the entity e using only the document D and the explicit entity mentions \mathcal{M}_e . We consider using external knowledge, such as information about entities from knowledge bases, to be outside the scope and leave integration of such knowledge for future work.

4 Methods

Pre-trained Language Models (PLMs) have shown a remarkable ability to encode syntactic and semantic knowledge in their parameters (Tenney et al., 2018, 2019) that can be leveraged when fine-tuned on downstream natural language understanding (NLU) tasks. We postulate that PLMs can be harnessed to help in entity salience detection, a target-based document-level NLU task. In this section, we present an architecture based on the cross-encoder setup adapted to the task of entity salience detection.

4.1 Cross-encoder

Encoding Given a document D and a target entity e , which is mentioned in the document, we concatenate the target entity’s name and the document using a special [SEP] token. We then encode the text using a Transformer-based pre-trained encoder. Figure 2 shows the graphical representation of the cross-encoder model. This setup allows the model to have deep cross attention between the target entity and the entire document. Note that we use special marker tokens [BEGIN_ENTITY] and [END_ENTITY] around each mentions $m \in \mathcal{M}_e$ of entity e in document D .

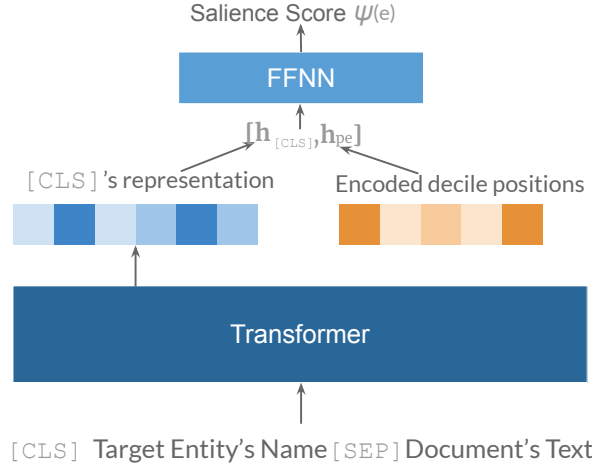


Figure 2: Graphical representation of the cross-encoder architecture with decile position encoding.

Position Encoding We compute the decile positions for each entity mention ($m \in \mathcal{M}_e$) in the document D by taking a positional index $p_m \in \{0, 1, \dots, 9\}$, indicating which part of document the mention belongs to if the document is partitioned into 10 equal chunks. Depending on the number and positions of the mentions, the vector can contain multiple non-zero values in the p vector. For example, if an entity e has 1 mention in the first decile, 2 in the second decile, and 1 mention in the fifth decile, then the input to the positional encoder would be $p_m = [1, 1, 0, 0, 1, 0, 0, 0, 0, 0]$. Note that we do not capture the number of mentions in each decile in p_m . To obtain positional embeddings, we use an embedding layer that maps positional indices to a dense vector of dimension d_{model} , formally $\mathbf{h}_{pe}(m) = \text{Embedding}(p_m)$.

Scoring The output representation of the [CLS] token is concatenated with the mean position embedding vector \mathbf{h}_{pe} and fed to a scorer module that produces a salience score $\psi(e) \in [0, 1]$ for entity e . The salience scorer is a feed-forward network with a sigmoid scoring function head. Formally,

$$\psi(e) = \sigma(\text{FFN}(\mathbf{h}_{[\text{CLS}]} \parallel \mathbf{h}_{pe}))$$

4.2 Optimization

We fine-tune the model described above by minimizing the binary cross entropy loss that is calculated using the ground truth binary salience labels and the predicted salience score $\psi(e)$.

5 Datasets

In this section, we describe our entity salience benchmark, which consists of four datasets: two

Dataset	NYT-Saliency	WN-Saliency	SEL	EntSUM
# Docs	110,463	6,956	365	693
Doc Length (avg chars)	5,079	2,106	1,660	4,995
# Unique entities	179,341	23,205	6,779	7,854
# Mentions	4,405,066	145,081	19,729	20,784
% Salient entities	14%	27%	10%	39%
Ground-truth	Abstract Alignment	Category Alignment	Human	Human

Table 1: Summary statistics and label collection methods for the datasets used in our experiments.

datasets were curated using semi-automated methods and two used human annotations. We provide summary statistics of these datasets and label collection methods in Table 1.

NYT-Saliency This dataset is introduced in (Dunietz and Gillick, 2014) and is the largest dataset to date for entity saliency detection. The dataset is curated with an assumption that salient entities are mentioned in the abstract of a news article in the NYT Corpus (Sandhaus, 2008). Entities and their mentions are identified using a classical NLP pipeline involving POS tagging, dependency parsing, and noun phrase extraction. Despite being large-scale, the automatic dataset creation process could introduce noise as corroborated by moderate agreement numbers with human annotators on a subset of the data. The dataset contains a binary saliency label for each entity.

WN-Saliency Introduced in (Wu et al., 2020a), this is another automatically curated dataset consisting of Wikinews articles. These are annotated with Wikinews categories by their authors. WN-Saliency identifies salient entities by using the hypothesis that an entity is salient if the Wikinews category that corresponds to the entity is also labeled as a category of the article. Similar to NYT-Saliency, this dataset has binary saliency labels.

SEL This is another dataset based on Wikinews released by (Trani et al., 2018). However, unlike WN-Saliency, this dataset is human annotated, where multiple human annotators ranked the saliency of entities into one of four categories. To conform with the binary labels of the other datasets, we map the 4 categories into binary labels of $\{0, 1\}$ by mapping the bottom two classes to not salient and the top two classes to salient.

EntSUM This dataset was introduced in (Maddala et al., 2022). To construct this dataset, a randomly selected set of entities spanning a subset of 693 articles from the NYT corpus were assigned

saliency labels by human annotators on a four-point scale, ranging between $[0, 3]$. For each document entity pair, two independent annotations were collected, which were increased up to 5 in case of disagreements. If the average annotation score is greater than 1.5 for an entity, it is assigned a positive saliency label.

5.1 Data Enrichment with Inferred Mentions

Except for EntSUM, the datasets do not have explicit entity mention offsets as annotations, which are necessary for many feature-based approaches and to compute positional embeddings. While SEL contains only the mention surface texts per entity, NYT-Saliency and WN-Saliency only provide the start and end character indices (aka mention offsets) of the very first mention of an entity. To this end, we infer additional mentions of an entity within the text using a combination of Flair NER (Akbi et al., 2019) and pattern matching.

For SEL, since the mentions are available, we use a pattern matching approach to match the surface text of the mentions to infer mention offsets. For NYT-Saliency and WN-Saliency, we first use Flair NER to identify mentions of named entities in the text. We attempt to match these mentions to the first mention of each entity in the document provided in the respective datasets. Since the surface text of other mentions may differ from the first mention, we additionally use the overlap between a mention’s surface text and the entity name as a candidate mention for that entity. Applying this approach, we infer additional mentions of an entity in the text and their offsets. While this process could introduce some noise, the overall quality of the datasets are enhanced through this process.

6 Experiments

We experiment on our entity saliency benchmark with our proposed PLM-based method, other ML and heuristic-based approaches used in past research, as well as an instruction-tuned PLM.

6.1 Data Splits

Prior works (Dunietz and Gillick, 2014; Trani et al., 2018; Wu et al., 2020a) use inconsistent (or not reported) train/validation/test splits. NYT-Salience and WN-Salience datasets are provided with train/test splits (but no validation), whereas SEL dataset is provided without any splits. This makes it hard to benchmark previous works with a fair comparison across models. To overcome this issue, we do a temporal split of NYT-Salience’s and WN-Salience’s original training sets into a new train/validation sets based on the publication time of the news stories, which provides a more realistic testing setup (Huang and Paul, 2018; Rijhwani and Preotiuc-Pietro, 2020). We also perform a temporal split of SEL and EntSUM datasets into train/validation/test sets. Further details about the dataset splits are provided in Appendix A.

6.2 Baselines

First, we list all methods used in past research, for which we report the results from the original papers.

- *First Sentence*. Classifies an entity as salient if it appears in the first sentence of the document’s body; used in both (Dunietz and Gillick, 2014) and (Wu et al., 2020a).
- *Position & Frequency* (Dunietz and Gillick, 2014). Feeds the first sentence index and the frequency features of an entity into a logistic regression model.
- *All Features* (Dunietz and Gillick, 2014). Uses a series of features based on position, frequency, and PageRank signals fed into a logistic regression model.
- *SEL* (Trani et al., 2018). Uses a combination of features based on position, frequency, and Wikipedia graph statistics fed into a Gradient Boosted Decision Tree algorithm implemented in sklearn (Pedregosa et al., 2011).
- *SWAT* (Ponza et al., 2019). Uses a set of features similar to the SEL Method described above, with the addition of features based on entity embeddings. All features are fed into a Gradient Boosted Decision Tree algorithm implemented in XGBoost (Chen et al., 2015).
- *Positional Feature* (Wu et al., 2020a). Uses the index of the first sentence in which the entity is mentioned as a feature in a logistic regression model. This method provides best results on the WN Salience dataset in (Wu et al., 2020a).

Next, we re-implement a set of common methods based on the above baselines in order to be able to test them on all four datasets. This ensures the evaluation is performed on the same experimental setup.

- *Positional Headline*. Classifies an entity as salient whether it appears in the headline of the input document.
- *Positional Headline & Lead*. Classifies an entity as salient if it appears in the headline of the document or in the first sentence (lead sentence) of the document.
- *Entity Frequency*. Classifies an entity as salient if they are more frequent than a given value. For each dataset, we calculated different thresholds and reported the best results. Thresholds can be found in the Appendix.
- *Features & GBDT*. This method uses the most common features from past works (Dunietz and Gillick, 2014; Wu et al., 2020a; Trani et al., 2018; Ponza et al., 2019) — i.e., entity’s first sentence index, and entity frequency — and feeds them into a GBDT model implemented using LightGBM (Ke et al., 2017).
- *SEL GBDT*. Follows the method from (Trani et al., 2018) and uses sklearn’s GBDT (Pedregosa et al., 2011) to train a model on the features provided with the SEL dataset.
- *Target entity masking*. This method feeds the input to a Transformer-based encoder (RoBERTa-base) with the target entity mentions represented through a special mask token. The salience prediction is obtained by mean pooling the mask token representations and passing this through a feed-forward network.
- *Zero-shot prompting*. We test instruction-tuned LLMs using zero-shot prompting. The prompt introduces the task description, followed by the input text and a target entity, and it asks a yes/no question. It expects the model to generate either ‘Yes’ or ‘No’ as an answer. The LLMs, already instruction-tuned on a large collection of NLU tasks, attempt to provide an answer based on the prompt, input text, and target entity. This family of models has been demonstrated to be robust and versatile on multiple benchmarks (Chung et al., 2022). We use FLan-UL2 (20B) (Tay et al., 2023) and LLaMa 2-Chat (7B) (Touvron et al., 2023) for evaluation.

Source	Type	Method	NYT-Saliency			WN-Saliency		
			P	R	F1	P	R	F1
(Dunietz and Gillick, 2014)	Heuristic	First Sentence	59.5	37.8	46.2	–	–	–
(Dunietz and Gillick, 2014)	ML	Position & Frequency	59.3	61.3	60.3	–	–	–
(Dunietz and Gillick, 2014)	ML	All Features	60.5	63.5	62.0	–	–	–
(Ponza et al., 2019)	ML	SWAT	62.4	66.0	64.1	–	–	–
(Wu et al., 2020a)	Heuristic	First Sentence	56.0	41.0	47.3	47.9	53.2	50.4
(Wu et al., 2020a)	ML	Positional Feature	19.0	41.3	26.0	29.1	78.9	42.5
(Wu et al., 2020a)	ML	Features & GBDT	39.2	59.7	47.3	29.2	48.1	36.3
Our Implementations	Heuristic	Positional Headline	57.5	42.0	48.5	46.1	51.5	48.7
	Heuristic	Positional Headline & Lead	49.8	55.4	52.5	41.0	60.0	48.7
	Heuristic	Entity Frequency	53.7	53.3	53.6	37.3	61.9	46.6
	ML	Features & GBDT	61.0	57.4	59.2	46.2	53.3	49.5
	PLM (RoBERTa)	Target Entity Masking	64.6	50.2	56.5	57.0	65.4	60.9
Our Models	PLM (RoBERTa)	cross-encoder	75.9	87.1	81.1	71.8	73.6	72.7
	PLM (DeBERTa)	cross-encoder	77.5	87.4	82.1	71.5	78.3	74.8
	PLM (RoBERTa)	cross-encoder w/ position emb.	78.7	84.2	81.4	71.2	76.7	73.8
	PLM (DeBERTa)	cross-encoder w/ position emb.	75.9	88.4	81.7	73.3	76.1	74.7

Table 2: Results on the NYT-Saliency and WN-Saliency datasets. The ground-truth of these datasets was generated via abstract/category alignment. The top section presents results as originally reported in the source papers.

Source	Type	Method	SEL			EntSUM		
			P	R	F1	P	R	F1
(Trani et al., 2018)	ML	SEL (w/ 5-fold cross val.)	50.0	61.0	52.0	–	–	–
(Ponza et al., 2019)	ML	SWAT (w/ 5-fold cross val.)	58.0	64.9	61.2	–	–	–
Our Implementations	Heuristic	Positional Headline	26.6	78.4	39.7	60.7	18.5	28.4
	Heuristic	Positional Headline & Lead	22.1	87.1	35.3	51.2	31.6	39.1
	Heuristic	Entity Frequency	13.5	57.8	21.9	48.4	54.0	51.0
	ML	Features & GBDT	26.6	78.4	39.7	60.7	52.0	56.0
	ML	SEL GBDT	71.1	47.8	57.1	–	–	–
	PLM (RoBERTa)	Target Entity Masking	36.3	13.8	20.0	63.0	41.7	50.2
Our Models	PLM (RoBERTa)	cross-encoder	51.6	73.6	60.6	65.5	60.6	63.0
	PLM (DeBERTa)	cross-encoder	64.1	73.6	68.5	64.9	59.2	61.9
	PLM (RoBERTa)	cross-encoder w/ position emb.	63.0	69.9	66.3	67.5	57.0	61.8
	PLM (DeBERTa)	cross-encoder w/ position emb.	67.3	62.4	64.7	72.1	51.5	60.1

Table 3: Results on the SEL and EntSUM datasets. The ground-truth of these datasets was generated via human annotation. The top section presents results as originally reported in the source papers.

6.3 Experimental Setup

We use RoBERTa-base (Liu et al., 2019) and DeBERTa-v3-base (He et al., 2023) as the base PLM for experiments. For each of these base models, we train both a cross-encoder model and a cross-encoder model augmented with decile positional embeddings.

For training our proposed models, we use AdamW (Loshchilov and Hutter, 2019) as the optimizer. We perform a hyperparameter search for learning rate using the following set of values: $\{0.001, 0.0005, 0.0002, 0.0001, 0.00005\}$. We train our models for a maximum of 10 epochs with early stopping based on the validation set performance. We pick the best performing model checkpoints for each dataset based on the perfor-

mance on the validation set. In Tables 2 and 3, we report the performance of our models and the baselines using the standard classification metrics (i.e., Precision, Recall, and F1) on the positive (salient) class, following previous research on entity saliency.

For training and inference of each Transformer-based model, we use a single NVIDIA V100 GPU with 32GB GPU memory, 4 CPUs, and 128 GB of main memory.

6.4 Results

In Tables 2 and 3, we present the experimental results of the baselines and our proposed models on the four datasets described in Section 5.

Comparison with feature-based methods. We

observe that the cross-encoder model significantly outperforms all baseline models in F1 score. It also yields better precision compared to the baselines for three of the four datasets. Only for the SEL dataset does the SEL GBDT model trained on publicly available pre-computed features produce a model with better precision than the cross-encoder. We observe that adding the decile positional embedding with cross-encoder improves the precision across all datasets, but also degrades the recall in every dataset except NYT-Salience.

The Target Entity Masking approach, which also leverages contextual information with a transformer-based model yields mixed results. Overall, the model is able to obtain better precision than the feature-based models for all datasets except SEL, but the model suffers from poor recall across all datasets, resulting in significantly worse F1 scores especially when compared to cross-encoder models.

Our re-implementation of positional methods and GBDT methods are consistent with the performance reported in prior works. The variance in numbers can be attributed to the enrichment of datasets with inferred mentions (Section 5.1) and the explicit train/dev/test data split used in our experiments (Section 6.1).

6.5 Zero-shot prompting of large language models

We formulate the problem of salience detection with zero-shot prompting as follows: given a definition of entity salience task and document text, we ask the model to generate a "yes" or a "no" if a particular entity is salient or not. We experimented with two open source models (Flan-UL2 (20B) and LLaMa 2-Chat (7B)) available on Hugging Face¹, and present the results in Table 4. To the best of our knowledge, this is the first evaluation of zero-shot prompting of instruction-tuned models for the entity salience detection task. We observe that the LLaMa 2-Chat model with 7 billion parameters fails to yield any meaningful results as it produces only positive labels for all data points (hence we observe 100% recall). The Flan-UL2 model is able to generate both positive and negative labels. However, the precision remains too low across datasets. We further discuss causes for this performance in the Appendix (Section C), along with the implementation details. Overall, these ex-

periments suggest that entity salience detection is a unique task that is not similar to any other tasks these two models are instruction tuned on.

7 Analysis

In this section, we perform an analysis of model predictions in order to gain more insights into model behavior and understand potential avenues for further improvement. We thus break down performance by different factors including: the importance of inferring all entity mentions, the position of the first entity mention, and entity mention frequency.

7.1 Impact of Inferred Mentions

In Section 5.1, we inferred additional mentions of an entity for the NYT-Salience and WN-Salience datasets. We compare the performance of our best model that leverages multiple mentions of an entity to its version trained with only the first mentions of entities in a document. The specific input formats for this experiment are presented in Appendix B. The results in Table 5 show that doing so consistently improves the performance of our models across all datasets. In particular, for the largest dataset, NYT-Salience, our model achieves a substantial gain of 27.3 F1 points. This experiment showcases the importance of augmenting our datasets with additional mentions and the importance of explicitly modelling contextual information present around all entity mentions.

7.2 Stratified Analysis on First Mention Position

We compare our cross-encoder models against the Features & GBDT model, our re-implemented baseline that relies on the most popular features used in prior works (Dunietz and Gillick, 2014; Wu et al., 2020a; Trani et al., 2018). As shown in the results from Tables 2 and 3, among other features, positional features are most informative for salience. Intuitively, if an entity is mentioned in the headline or in the first sentence of a news article, there is high probability of that entity being salient.

Figure 3 shows that all models perform well when the first mention falls in the headline or the first sentence of the document. We notice that the cross-encoder models constantly outperform the Features & GBDT model and the largest gains are observed in the SEL and WN-Salience datasets. This observation indicates that the cross-encoder models are


¹www.huggingface.com

Model	NYT-Saliency			WN-Saliency			SEL			EntSUM		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Cross-encoder (DeBERTa)	77.5	87.4	82.1	71.5	78.3	74.8	64.1	73.6	68.5	64.9	59.2	61.9
Flan-UL2	31.1	72.4	43.5	30.7	90.1	45.9	16.7	98.3	28.5	27.6	83.6	41.5
LLaMa 2-Chat	14.6	100.0	25.4	27.1	100.0	42.6	9.49	100.0	17.3	19.2	100.0	32.2


Table 4: Performance comparison of cross-encoder model with zero-shot prompting of LLMs.

Model	NYT-Saliency			WN-Saliency			SEL			EntSUM		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Cross-encoder w/ first mention	54.2	57.5	55.8	69.6	80.4	74.6	59.8	76.1	67.0	69.1	53.2	60.2
Cross-encoder w/ all mentions	77.5	87.4	82.1	71.5	78.3	74.8	64.1	73.6	68.5	64.9	59.2	61.9

Table 5: Performance comparison of cross-encoder models with only the first mention vs. all inferred mentions.



(a) Performance with respect to the position of the mentions. There are no mentions outside of the context window for NYT.



(b) Performance with respect to the frequency of the entities. The test split of SEL dataset does not contain any entity with more than 10 mentions in a document.

Figure 3: Stratified analysis across models and datasets.

able to use the context to identify that mentions that occur in the headline or the first parts of the document are often salient without explicitly using this information as a feature.

We also investigate the performance of the models when the first mention falls inside or outside the context window of the PLM (here, 512 tokens). When mentions fall inside the context window, we observe that the cross-encoder models consistently outperform the Features & GBDT model. When the mention falls outside the context window, the model predictions become close to random, which is expected, as the model does not have immediate contextual information around the mention. Using models that can deal with longer inputs would be a promising direction for improvement for these samples (Beltagy et al., 2020). Interestingly, for WN-Salience, the Features & GBDT model also performs considerably worse outside the first 512 tokens.

7.3 Stratified Analysis on Mention Frequency

Similar to mention position analysis, we compare our cross-encoder models against the Features & GBDT model, which uses mention frequency as one of its input features. Figure 3 shows how the cross-encoder models and Features & GBDT compare with varying frequency of entity mentions.

For salient entities with single mentions, the cross-encoder model performs significantly better than the Features & GBDT model. In particular, for the NYT-Salience dataset, the Features & GBDT model fails to predict any of the single mention entities as salient. This observation indicates that the cross-encoder models do not simply model the mention frequency, but potentially leverage other contextual information to determine the salience of entities with a single mention.

The performance of the Features & GBDT model improves with more mentions per entity. In fact, for the frequency range of 6-10 mentions per entity, the Features & GBDT model performs better than the cross-encoder models for EntSUM and SEL datasets. This observation indicates the over-reliance of the Features & GBDT model on mention frequency to determine salience, but also that the cross-encoder cannot fully use this heuristic.

8 Conclusion

This paper aims to leverage the semantic knowledge encoded in pre-trained language models for

entity salience detection. We propose the cross-encoder method based on Transformer-based PLMs with positional representations and compare its performance to several ML-based methods, heuristic methods, and instruction-tuned LLMs across four different datasets, two human-annotated and two automatically curated. Across all our experiments, the cross-encoder model based on pre-trained language models outperforms all other methods, often with double digit gains in F-1 score. Analyses of model behavior illustrate the important effects of mention frequency, mention position, and document length on performance, highlighting areas of future work.

9 Limitations

We only studied salience in English-language documents, but our methods are applicable to other languages directly as long as a pre-trained language model covering the target language is available.

We use entity mentions as annotated in our data or inferred through entity recognition and entity resolution for inference in some of the methods. This information may not be available at inference time in all applications.

The experiments with LLMs are limited to zero-shot prompts. We did not experiment with instruction tuning which could potentially help the model learn the salience detection task.

Finally, we do not use external knowledge about entities and their relationships in modelling, which was shown to marginally improve results in past studies (Dunietz and Gillick, 2014; Trani et al., 2018; Ponza et al., 2019). We consider this out of the scope of our analysis and a viable direction of future work.

10 Ethics Statement

We use publicly available datasets intended for the task of entity salience detection. The datasets and pre-trained models we used have permissive licenses allowing for research use. We do not envision any potential risks associated with the task discussed in this paper.

References

Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. FLAIR: An easy-to-use framework for state-of-the-art NLP. In *NAACL 2019, 2019 Annual Conference of the North American Chapter of the Association for*

- Computational Linguistics (Demonstrations)*, pages 54–59.
- Zakariae Alami Merrouni, Bouchra Frikh, and Brahim Ouhbi. 2020. Automatic keyphrase extraction: a survey and trends. *Journal of Intelligent Information Systems*, 54(2):391–424.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(null):993–1022.
- Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, et al. 2015. XGBoost: Extreme Gradient Boosting. *R package version 0.4-2*, 1(4):1–4.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Leon Derczynski, Kalina Bontcheva, and Ian Roberts. 2016. [Broad Twitter corpus: A diverse named entity recognition resource](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1169–1179, Osaka, Japan. The COLING 2016 Organizing Committee.
- Milan Dojchinovski, Dinesh Reddy, Tomáš Kliegr, Tomáš Vitvar, and Harald Sack. 2016. [Crowdsourced corpus with entity salience annotations](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3307–3311, Portorož, Slovenia. European Language Resources Association (ELRA).
- Jesse Dunietz and Daniel Gillick. 2014. [A new entity salience task with millions of training examples](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 205–209, Gothenburg, Sweden. Association for Computational Linguistics.
- Günes Erkan and Dragomir R. Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *J. Artif. Int. Res.*, 22(1):457–479.
- Michael Gamon, Tae Yano, Xinying Song, Johnson Apacible, and Patrick Pantel. 2013. [Understanding document aboutness step one: Identifying salient entities](#). Technical Report MSR-TR-2013-73.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [DeBERTav3: Improving deBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations*.
- Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenau, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. [Robust disambiguation of named entities in text](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 782–792, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Ella Hofmann-Coyle, Mayank Kulkarni, Lingjue Xie, Mounica Maddela, and Daniel Preotiuc-Pietro. 2022. [Extractive entity-centric summarization as sentence selection using bi-encoders](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 326–333, Online only. Association for Computational Linguistics.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. [OntoNotes: The 90% solution](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Xiaolei Huang and Michael J. Paul. 2018. [Examining temporality in document classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 694–699, Melbourne, Australia. Association for Computational Linguistics.
- Anette Hulth. 2003. [Improved automatic keyword extraction given more linguistic knowledge](#). In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 216–223.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. [End-to-end neural entity linking](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 519–529, Brussels, Belgium. 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*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Mounica Maddela, Mayank Kulkarni, and Daniel Preotiuc-Pietro. 2022. [EntSUM: A data set for entity-centric extractive summarization](#). In *Proceedings of*

- the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 3355–3366, Dublin, Ireland. Association for Computational Linguistics.
- Edgar Meij, Wouter Weerkamp, and Maarten de Rijke. 2012. [Adding semantics to microblog posts](#). In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining, WSDM '12*, page 563–572, New York, NY, USA. Association for Computing Machinery.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Dat Ba Nguyen, Johannes Hoffart, Martin Theobald, and Gerhard Weikum. 2014. [Aida-light: High-throughput named-entity disambiguation](#). In *Proceedings of the Workshop on Linked Data on the Web co-located with the 23rd International World Wide Web Conference (WWW 2014), Seoul, Korea, April 8, 2014*, volume 1184 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Marco Ponza, Diego Ceccarelli, Paolo Ferragina, Edgar Meij, and Sambhav Kothari. 2021. Contextualizing trending entities in news stories. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 346–354.
- Marco Ponza, Paolo Ferragina, and Francesco Piccinno. 2019. [Swat: A system for detecting salient wikipedia entities in texts](#). *Computational Intelligence*, 35(4):858–890.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. [Towards robust linguistic analysis using OntoNotes](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152, Sofia, Bulgaria. Association for Computational Linguistics.
- Shruti Rijhwani and Daniel Preotiu-Pietro. 2020. [Temporally-informed analysis of named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7605–7617, Online. Association for Computational Linguistics.
- Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- Benjamin Strauss, Bethany Toma, Alan Ritter, Marie-Catherine De Marneffe, and Wei Xu. 2016. Results of the w-nut 2016 named entity recognition shared task. In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 138–144.
- Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Siamak Shakeri, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. 2023. [Ul2: Unifying language learning paradigms](#).
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2018. What do you learn from context? probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Salvatore Trani, Claudio Lucchese, Raffaele Prego, David E. Losada, Diego Ceccarelli, and Salvatore Orlando. 2018. [Sel: A unified algorithm for salient entity linking](#). *Computational Intelligence*, 34(1):2–29.

- Chuan Wu, Evangelos Kanoulas, Maarten de Rijke, and Wei Lu. 2020a. Wn-salienc: A corpus of news articles with entity salience annotations. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 2095–2102.
- Chuan Wu, Evangelos Kanoulas, and Maarten de Rijke. 2020b. [Learning entity-centric document representations using an entity facet topic model](#). *Inf. Process. Manage.*, 57(3).
- Chenyan Xiong, Zhengzhong Liu, Jamie Callan, and Tie-Yan Liu. 2018. Towards better text understanding and retrieval through kernel entity salience modeling. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 575–584.
- Lingyun Zhao, Lin Li, Xinhao Zheng, and Jianwei Zhang. 2021. A bert based sentiment analysis and key entity detection approach for online financial texts. In *2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, pages 1233–1238. IEEE.

Appendix

A Details of Dataset Splits

Table 6 contains the train, dev, and test splits of each of the datasets after applying a temporal splitting strategy described in Section 6.1. These splits are used for model training and evaluation.

B Input Format for Experiments

As described in Section 4.1, we add special marker tokens around each mention of the target entity (i.e., the entity for which the model needs to predict the salience label.). In the following, we provide an example:

Text
Musk completes \$44 billion Twitter deal. Elon Musk, the world's ...
Model Input
[CLS] Elon Musk [SEP] [BEGIN_ENTITY] Musk [END_ENTITY] completes \$44 billion Twitter deal. [BEGIN_ENTITY] Elon Musk [END_ENTITY], the world's ...

For the experiment with first mention reported in Section 7.1, only the first mention is bounded by special marker tokens as shown in the following example:

Model Input
[CLS] Elon Musk [SEP] [BEGIN_ENTITY] Musk [END_ENTITY] completes \$44 billion Twitter deal. Elon Musk, the world's ...

Note that the second mention of Elon Musk is not bounded by marker tokens.

C Implementation details of zero-shot prompting of LLMs

Figure 4 and Figure 5 show the prompts we used for the LLaMa 2-Chat (7B) and FLan-UL2 (20B) models respectively. Table 7 lists the generation parameters. We speculate the following causes for the relatively lower precision obtained using this method:

- The instruction defines the salience task definition, but doesn't provide any reference ex-

amples (few-shot prompting) to align with the definition of salience. This leads to the model identifying an entity as salient based on its frequency in the document. However, creating a few-shot prompt is challenging as we need to limit the maximum input length of the prompt to prevent out-of-memory issues.

- We truncate the document text so that the entire prompt is 2048 tokens or less, thus throwing away any potential information present towards the end of a long document.

```
<s> [INST]
«SYS» The salience of an entity provides information about the importance or centrality of that entity to the entire document text. In the following, given an Entity and a Text, you need to answer 'Yes' if the Text document is about that Entity and 'No' if the Text is not about that Entity. «/SYS»
Is Entity: entity salient in Text: text
[/INST]
```

Figure 4: Instruction for zero-shot prompting of LLaMa 2-Chat model.

```
### Instruction ###
The salience of an entity provides information about the importance or centrality of that entity to the entire document text. In the following, given an Entity and a Text, you need to answer 'Yes' if the Text document is about that Entity and 'No' if the Text is not about that Entity.
Text: text
Entity: entity
Question: Is the above Entity salient in the above Text? Please answer Yes or No.
Answer:
```

Figure 5: Instruction for zero-shot prompting of FLan-UL2 model.

D Thresholds for Entity Frequency baseline

Figure 6 shows the performance of the Entity Frequency baseline by varying the minimum number of times an entity has to occur in the input document to be classified as salient.

Dataset	# Doc-Entity pairs	Train	Validation	Test
NYT-Saliency	1,910,214	1,342,092	405,335	162,787
WN-Saliency	62,537	41,625	11,902	9,009
SEL	12,257	6,106	2,400	3,751
EntSUM	9,934	5,206	1,861	2,867

Table 6: Document-Entity pairs in train, validation, and test splits after applying temporal splitting.

Generation parameter	Value
top_k	0
top_p	0
temperature	0
max_new_tokens	1

Table 7: Parameters for generating a saliency label with zero-shot prompt.




Figure 6: Performance of the Entity Frequency baseline over different thresholds.

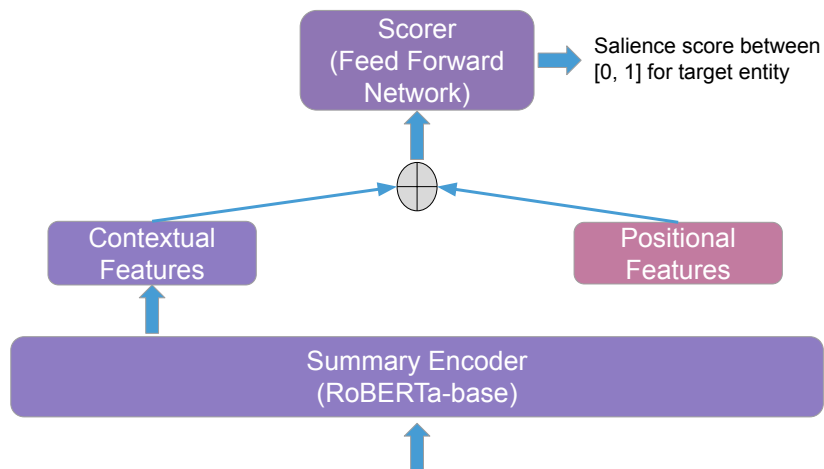


Figure 7: Schematic diagram of the Target Entity Masking model architecture.