

Comparative Analysis of Named Entity Recognition in the Dungeons and Dragons Domain

Gayashan Weerasundara and Nisansa de Silva

Department of Computer Science & Engineering

University of Moratuwa, Sri Lanka

{gayashan.22, NisansaDds}@cse.mrt.ac.lk

Abstract

Many NLP tasks, although well-resolved for general English, face challenges in specific domains like fantasy literature. This is evident in Named Entity Recognition (NER), which detects and categorizes entities in text. We analyzed 10 NER models on 7 Dungeons and Dragons (D&D) adventure books to assess domain-specific performance. Using open-source Large Language Models, we annotated named entities in these books and evaluated each model's precision. Our findings indicate that, without modifications, Flair, Trankit, and Spacy outperform others in identifying named entities in the D&D context.

1 Introduction

Named Entity Recognition (NER) targets the identification and classification of textual entities, such as names and locations. In the diverse and intricate vocabulary of fantasy literature, like that of Dungeons and Dragons (D&D), NER becomes challenging (Zagal and Deterding, 2018). D&D, a prominent fantasy literature domain, spans content for its namesake tabletop game (Peiris and de Silva, 2022, 2023; Zhou et al., 2022). These narratives inhabit fictional realms like Forgotten Realms and Dragonlance, bursting with characters, locations, and objects (Gygax and Arneson, 1974).

NER's utility in fantasy literature is vast: from extracting information and summarizing text to character analysis and plot creation. However, conventional NER models, primarily trained on standard datasets like CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003) or OntoNotes 5.0 (Weischedel et al., 2013), might falter on fantasy texts due to their unique linguistic attributes. Recognizing the need for domain-specific adaptation, other specialized areas such as law (Sugathadasa et al., 2017), medicine (de Silva et al.,

2017), and the dynamic landscape of social media (de Silva and Dou, 2021) have already seen research emphasizing it. Large models, as Yao et al. (2021) points out, can face domain adaptation challenges, stressing the need for evaluating NER models specifically on fantasy content.

Fantasy NER has potential, especially with advancements in image generation. A notable application might involve an image generation model leveraging NER tags to derive prompts and subsequently produce contextually relevant images.

Our study contrasts 10 NER models across seven D&D books, each averaging 118,000 words. Manual annotations of entities were made and juxtaposed against model outputs. Through precision assessments and named entity distribution analyses, we glean insights into model performances in the fantasy domain. Our key contributions include:

- A pioneering, comprehensive NER model evaluation on fantasy content.
- An annotated D&D book dataset for NER studies.
- A deep dive into varied NER models' strengths and pitfalls in the fantasy realm.
- Discussions on NER's role and prospects in fantasy literature.

Following this, Section 2 delves into related NER and fantasy literature works. Section 3 details our data and annotation process, while Section 4 unveils our methods and findings. Sections 5 and 6 respectively discuss insights and conclude our research, and Section 7 outlines potential future endeavors.

2 Related Works

NER has seen the development of various models like rule-based systems, statistical models, neural

networks, and transformer-based models (Seo et al., 2021; Liu et al., 2022; Krasnov et al., 2021). Although they’ve been trained on standard datasets, these don’t encompass the complexities found in domains like fantasy literature, which poses challenges due to invented names, variable spellings, entity ambiguity, and limited resources.

We introduced a novel annotated dataset of D&D books for NER and evaluated 10 NER models, including XLM-RoBERTa (Conneau et al., 2019), StanfordDeID (Chambon et al., 2023), ELECTRA (Clark et al., 2020), and others.

Other studies have compared the performance of NER models on different types of texts and languages. For example, Wang et al. (2021) compared Spacy, Flair, m-BERT, and camemBERT on anonymizing French commercial legal cases. They found that camemBERT performed the best overall, followed by Flair and m-BERT. SpaCy had the lowest scores but also the fastest prediction time. (Benesty, 2019) compared spaCy, Flair, and Stanford Core NLP on anonymizing English court decisions. They found that Flair had the highest scores, followed by Stanford Core NLP and spaCy. (Shelar et al., 2020) compared rule-based, CRF-based, and BERT-based techniques for NER on text data. They found that BERT-based technique had the highest accuracy and recall, followed by CRF-based and rule-based techniques. (Naseer et al., 2021) compared NLTK, spaCy, Stanford Core NLP, and BERT (Devlin et al., 2018) on extracting information from resumes. They found that BERT had the highest accuracy and F-measure, followed by spaCy, Stanford Core NLP, and NLTK.

These studies suggest that different NER models may have different strengths and weaknesses depending on the type, language, and domain of the text data. Our study aims to contribute to this understanding by providing the first systematic comparison of NER models on fantasy texts and analyzing their performance and characteristics.

3 Data Collection and Annotation

This section details the data sources and annotation process utilized for our named entity recognition (NER) task, a subtask of information extraction that classifies named entities in unstructured text into categories such as persons, organizations, and locations (Mohit, 2014).

We examined seven adventure books from the Dungeons and Dragons (D&D) realm, listed in

table 1. These books, primarily adventure-centric, were sourced from the official DnDBeyond site, the main publication hub for D&D by Wizards of the Coast.

Through a comprehensive analysis of these texts, we used their rich narratives and character dynamics to benchmark and assess various NER models in this intricate domain.

Book	Counts	
	Words	Topics
Lost mine of Phandelver (Baker and Perkins, 2014)	45947	29
Hoard of the Dragon Queen (Baur et al., 2014a)	74243	45
Rise of Tiamat (Baur et al., 2014b)	80065	48
Curse of Strahd (Perkins et al., 2016)	154519	62
Tomb of Annihilation (Perkins et al., 2017)	148605	35
Candlekeep Mysteries (Perkins et al., 2021)	141104	106
The Wild Beyond the Witchlight (Allan et al., 2021)	184135	60

Table 1: D&D adventure books

Each of our chosen books averages 118,000 words. The selection was driven by our familiarity with these tales and the broader D&D universe. Additionally, they span multiple genres, themes, and settings in the fantasy realm, offering a vast array of named entities for NER.

The source books were transformed into text and organized hierarchically into chapters, topics, and paragraphs. An example from "The Wild Beyond the Witchlight" is displayed in table 2.

We first manually perused the source books, marking named entities hierarchically by chapter, topic, and paragraph, recording only entity counts. Subsequently, we employed three state-of-the-art large language models: Bloom (Scao et al., 2022), OpenLLaMA (Geng and Liu, 2023), and Dolly (Databricks, 2023), to detect named entities in each book chapter. These models, trained on vast conversational data, can craft natural language responses, making them apt for the intricate language patterns in D&D texts, such as neologisms and metaphors.

After eliminating duplicates and pinpointing unique entities, we verified these results against our initial counts. The named entities identified by the three LLMs underwent a manual review for accuracy and consistency, adding crucial missed entities. Table 3 contrasts the named entity counts from each LLM, with recall metrics based on entities common across all models.

When annotating the resultant named entities we followed a set of annotation guidelines that define the entity types and the annotation rules for our NER task. The entity types that were used are:

Chapter	Topic	Paragraph	Word Count
Introduction: Into the Feywild	Adventure Summary	The main antagonists of this story are three hags...	131
		One of the many novelties of this adventure is that...	43
		The characters are drawn into the adventure by one of two adventure hooks.	31
		You choose... Chapter 1 describes the Witchlight Carnival...	40
Running the Adventure	
		The Monster Manual contains stat blocks for most of the creatures encountered in this...	72
		Spells and equipment mentioned in the adventure are described in the Players Handbook...	31

Table 2: Content hierarchy in a book

- **Person:** any named character or creature that can act as an agent, such as heroes, villains, allies, enemies, etc.
- **Organization:** any named group or faction that has a common goal or identity, such as guilds, cults, clans, etc.
- **Location:** any named place or region that has a geographical or spatial dimension, such as cities, dungeons, forests, etc.
- **Misc:** any named entity that do not belong to above mentioned categories. (This contain important information like Spells, Artifacts, Potions etc.)

The process of annotation is done through a script, where a paragraph segment is taken iteratively and fed into the LLMs with a template prompt.

Following Algorithm 1 is the pseudo-code for the process in identifying named entities:

As shown in above pseudocode, the algorithm 1 takes a set of books as input and outputs the named entities identified by the LLMs. The algorithm iterates over each book and divides it into segments. Each segment is further divided into paragraphs,

Input: Books;

Output: Named entities;

foreach *book* **do**

segments ← divideIntoSegments(book);

foreach *segment in segments* **do**

paragraphs ←

divideIntoParagraphs(segment);

foreach *paragraph in paragraphs*

do

foreach *LLM in LLMs* **do**

prompt ←

createPrompt(paragraph);

namedEntities ←

LLM(prompt);

processNamedEntities(namedEntities);

end

end

end

end

end

removeDuplicates(namedEntities);

Algorithm 1: Named Entity Recognition using Multiple LLMs

and each paragraph is iteratively fed into each of the LLMs with a prompt to identify named entities. The named entities identified by each LLM are then processed and saved. Finally, all named entities are checked for duplicates, and those duplicates are removed.

After named entities were recognized, they were then mapped in to json objects for storage as shown in Figure 1. Nesting of objects is done according to the hierarchy as mentioned in table 2. Each of the named entities were nested in an array of entities as entity objects with corresponding attributes as mentioned below.

4 Experimental Setup and Results

The experiment was conducted to identify how effective are the NER models when using them as off the shelf models in identifying named entities for a fantasy domain when there are no available corpora for fine tuning. For testing we used 10 different contemporary NER models.

Following table 5 shows the identified count of named entities for each categories of the adventure book Candlekeep Mysteries.

The testing approach for the NER models mirrors algorithm 1. Here, paragraphs of input text are fed into the models without specific prompts.

Book	Bloom		Dolly		OpenLLaMA		Total Unique Entities
	Count	Recall	Count	Recall	Count	Recall	
Lost Mine of Phandelver	21	0.47	32	0.73	40	0.91	44
Hoard of the Dragon Queen	58	0.89	62	0.95	60	0.92	65
Rise of Tiamat	54	0.88	57	0.93	53	0.87	61
Curse Of Strahd	92	0.90	96	0.94	101	0.99	102
Tomb of Annihilation	101	0.80	99	0.79	112	0.89	126
Candle keep Mysteries	60	0.87	61	0.88	64	0.93	69
The Wild Beyond Witch Light	66	0.84	67	0.85	71	0.89	79

Table 3: Result comparison between LLMs

Please identify and list all named entities in the following text using the BIO (beginning-inside-outside) scheme:		
”The traveling extravaganza known as the Witchlight Carnival visits your world once every eight years. You have a dim memory of sneaking into the carnival as a child without paying... ..pair of elves named Mister Witch and Mister Lightwere decidedly unhelpful.”		
B-Organization:	Witchlight	Carnival
I-Person:	Mister	Witch
I-Person:	Mister	Light

Table 4: Process of Annotation

Model	PER	LOC	ORG	MSC	All
XLM-RoBERTa (Conneau et al., 2019)	16	0	3	4	23
StanfordAIMI (Chambon et al., 2023)	0	0	1	18	19
ELECTRA (Clark et al., 2020)	10	0	1	10	21
WikiNeuRal (Tedeschi et al., 2021)	23	4	6	1	34
BERT (Devlin et al., 2018)	9	1	1	0	11
RoBERTaNER (Baptiste, 2022)	1	0	0	17	18
BERT-CRF (Souza et al., 2019)	12	0	0	0	12
Flair (Akbik et al., 2018)	28	14	6	4	54
Spacy (Honnibal and Montani, 2017)	21	11	7	18	57
Trankit (Nguyen et al., 2021)	25	15	2	2	44

Table 5: Statistics for the adventure book Candlekeep Mysteries. The NER tags are as follows, Person: PER, Location: LOC, Organization: ORG, and Miscellaneous: MSC

The resultant output is refined by filtering out corrupted values (e.g., ”Strahd Von Zarovich” might be mistakenly split into two distinct names) and redundant entries, before being transitioned into the JSON structure showcased in Figure 1.

During initial processing, NER models often produce numerous erroneous outputs. These arise from factors like incomplete word detection, mis-

```

JSON object
{
  "book": "Candlekeep_Mysteries",
  "chapter": 1,
  "text": "The_Book_of_Inner_Alchemy_is_one_of_Candlekeeps... ",
  "entities": [
    {
      ...
    },
    {
      "entity": "B-Location",
      "score": 0.9659823,
      "index": 8,
      "word": "Candlekeep",
      "start": 42,
      "end": 51
    },
  ]
}

```

Figure 1: sample format of the JSON output

segmentation of terms, or misinterpretation of special characters. Such discrepancies can be mitigated using string manipulations and by cross-referencing outputs with a pre-curated list of named entities.

Figure 2 displays entries that encountered corruption. These highlight instances where NER models incorrectly processed and extracted entities from the source material.

In the given example shown in Figure 2, the name ”Fembris Larlancer” is erroneously divided into two distinct words, ”Fembris L#” and ”rlancer”, as a result of corruption during the NER processing stage. This example underscores the challenges faced during the entity extraction pro-

```

JSON object
{
  "entities": [
    {
      "entity": "I-Person",
      "score": 0.5659823,
      "index": 308,
      "word": "Fembris_L#",
      "start": 1160,
      "end": 1162
    },
    {
      "entity": "I-Person",
      "score": 0.51227564,
      "index": 309,
      "word": "rlancer",
      "start": 1162,
      "end": 1164
    }
  ]
}

```

Figure 2: sample format of a corrupted JSON outputs

cess and the need for robust post-processing to ensure the accuracy and quality of the extracted entities.

After removing corrupted and eligible named entities, duplicate entries must be removed to do a proper comparison of performance between different models. For this tuples of words in adjacent positions were generated and compared. For example Mayor Lei Duvezin, Mayor Duvezin, Lei Duvezin and Duvezin all refers to the same entity with the label Person. In cases such as above tuple with most similarity matches will be retained as the named entity and duplicates will be removed.

To visualize the raw named entity identification potential of each model, a density plot was plotted with respect to count of identified named entities with NER models. Following Figure 3 shows the density of named entities recognized by each NER model. The hue represents the overlapping count ranges of named entities identified in each source book.

Without training, NLP frameworks like Trankit (Nguyen et al., 2021), Flair (Akbik et al., 2018), and Spacy (Honnibal and Montani, 2017) show a strong baseline in entity recognition.

Model precision is key in performance evaluation. This is gauged by comparing the true positive entities with actual named entities. This comparison can be visually represented for each model




Figure 3: Density plot for each model

across source books.

For a comprehensive model assessment across books, Kernel Density Estimation (KDE) is used. It's a non-parametric method estimating the probability density function (Terrell and Scott, 1992):

$$f(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right)$$

where:

- x_i are the data points
- K is the kernel function, which is typically a Gaussian function or a uniform function
- h is the bandwidth, which determines the width of the kernel function and controls the smoothness of the estimate
- n is the number of data points

KDE calculates $f(x)$ through a summed kernel function $K(u)$, anchored at data points x_i .

Figure 4 illustrates models' efficacy over seven source books. A gradient near 1 signifies optimal performance.

In DD, named characters, with their elaborate backstories, are central. Assessing a model's inclination to identify these characters over other entities is vital. This inclination can be visualized

by juxtaposing character counts with total entities, contrasted against real metrics. Figure 5 delineates the frequency of character identification across all source books. Meanwhile, Figure 5a and Figure 5b depict the distribution pertaining to models and books, respectively.

In the D&D landscape, named characters, renowned for their intricate histories, are paramount. Evaluating a model’s propensity to spot these characters in relation to other entities is imperative due to the significant role characters play in D&D narratives. This bias can be graphically represented by mapping character counts against all identified entities and contrasting them with authentic counts. By scrutinizing the named entity counts from diverse NER models and comparing them to true values, one can infer model behavior and efficacy. Figure 5 offers a glimpse into character recognition frequency for different models across sourcebooks, with Figure 5a and Figure 5b charting the distributions for models and books respectively.

From Figure 5b, we observe a consistent ratio between characters and other named entities across books. This consistency allows us to downplay book variability and focus on the insights from Figure 5a. Notably, NLP frameworks such as Spacy (Honnibal and Montani, 2017) and Flair (Akbi et al., 2018) exhibit more balanced frequency distributions, indicating a higher character identification ratio. Although this might be unfavorable in certain contexts, in this domain, aligning character identifications closely with overall named entity values signals optimal performance. This suggests Spacy and Flair perform exceptionally in an off-the-shelf setting.

Figure 6 showcases precision and recall metrics for each NER model. To determine recall, we derived the true positive count from average unique named entity counts, while the true count originated from LLM models, as outlined in table 3. For precision, false positives were ascertained from misidentified unique named entities on average.

The precision and recall values were averaged for each model across source books, and plotted to offer a concise visualization of each NER model’s performance.

Evidently, Flair and Spacy outshine other NER models in precision and performance, while Trankit (Nguyen et al., 2021) excels in recall relative to its precision.

5 Discussion

We undertook a Named Entity Recognition (NER) task on seven adventure books from the esteemed Dungeons and Dragons (D&D) series. Our methodology involved manual entity annotations in these books, which were subsequently verified against outputs from three leading language models: Bloom, OpenLLaMA, and Dolly.

Our annotation guidelines delineated entity types into categories like person, organization, location, and misc. Ten NER models were subsequently employed to gauge their efficacy in recognizing named entities within D&D. Among these, Flair, Trankit, and Spacy emerged superior, mirroring findings from past NER-centric studies. Conversely, StanfordDeID (Chambon et al., 2023) and RoBERTaNER (Baptiste, 2022) lagged in performance. A precision-centric analysis further reiterated the dominance of Flair, Trankit, and Spacy over their counterparts.

The findings imply that while generic models can decently handle NER tasks in specialized domains like D&D, performance inconsistencies exist across models. Employing annotation guidelines bolsters consistency in entity recognition across varied books and contexts. Moreover, incorporating large language models for automated annotations can significantly mitigate the manual intervention needed for comprehensive datasets, particularly in intricate domains such as D&D.

However, our study bears certain caveats. We refrained from fine-tuning the NER models specifically for D&D, so our findings are indicative of generic model capabilities and might not capture the full potential of domain-specific optimization. Our dataset, comprising just seven books, might not encompass the depth and breadth of D&D narratives. The exclusive focus on Wizards of the Coast publications could also inadvertently introduce stylistic biases. Finally, while our study zeroes in on D&D as a fantasy subset, our insights might not seamlessly extend to other literary domains with their unique nuances.

6 Conclusion

Our exploration illuminates the remarkable potential of harnessing off-the-shelf models for NER tasks within the D&D universe’s nuanced realm. Some models showcase an impressive baseline in entity recognition for this domain without extensive fine-tuning. However, there’s a compelling need for




Figure 4: Distribution plot for each model




Figure 5: Frequency plots with respect to models and adventure sourcebooks

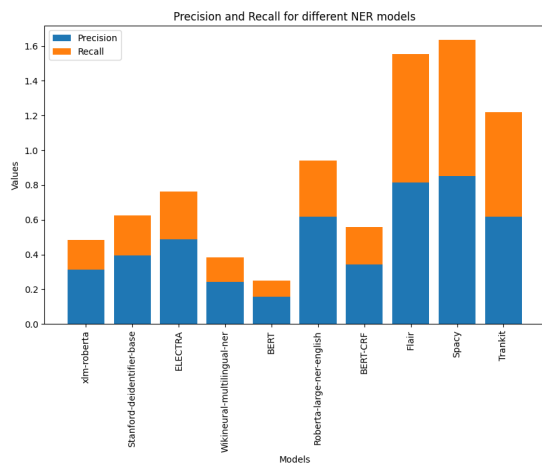


Figure 6: Precision graph for different NER models

continued research and refinement to tailor these models optimally for D&D’s unique intricacies.

Additionally, our research serves as a foundational resource for future inquiries. The dataset we’ve curated and our annotation guidelines stand as a benchmark for gauging the efficiency of future NER models or techniques. Consequently, our work not only reveals the current prowess of NER

models within the D&D context but also sets the stage for continued innovation at the confluence of fantasy literature and artificial intelligence.

7 Future Works

Based on our findings and limitations, we suggest some directions for future research. One direction is to fine-tune NER models on the D&D dataset and comparing their performance with off-the-shelf models. Additionally, other techniques such as transfer learning or domain adaptation could be explored to improve the performance of NER models in the D&D domain. Another direction is to use different data sources for NER in D&D, such as novels, comics, podcasts, or video games. A third direction is to apply different evaluation metrics for NER in D&D, such as F1-score, recall, accuracy, or error analysis. Finally other aspects of NER in D&D can also be explored, such as entity linking, coreference resolution, relation extraction, or sentiment analysis.

References

- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*, pages 1638–1649.
- Stacey Allan, Will Doyle, Ari Levitch, and Christopher Perkins. 2021. *The Wild Beyond the Witchlight*. Wizards of the Coast.
- R Baker and Christopher Perkins. 2014. Lost mine of phandelver. *Wizards of the Coast*.
- Jean Baptiste. 2022. roberta-large-ner-english: model fine-tuned from roberta-large for ner task. <https://huggingface.co/Jean-Baptiste/roberta-large-ner-english>. (Accessed on 05/10/2023).
- Wolfgang Baur, Steve Winter, and Alexander Winter. 2014a. *Hoard of the Dragon Queen*. Wizards of the Coast.
- Wolfgang Baur, Steve Winter, and Alexander Winter. 2014b. *Rise of Tiamat*. Wizards of the Coast.
- M Benesty. 2019. Ner algo benchmark: spacy, flair, m-bert and camembert on anonymizing french commercial legal cases. *Toward. Data Sci.*
- Pierre J Chambon, Christopher Wu, Jackson M Steinkamp, Jason Adleberg, Tessa S Cook, and Curtis P Langlotz. 2023. Automated deidentification of radiology reports combining transformer and hide in plain sight rule-based methods. *Journal of the American Medical Informatics Association*, 30(2):318–328.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Databricks. 2023. [Dolly: The first open source, commercially viable instruction-tuned llm](#). *Databricks Blog*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Xinyang Geng and Hao Liu. 2023. [Openllama: An open reproduction of llama](#).
- Gary Gygax and Dave Arneson. 1974. *dungeons & dragons*, volume 19. Tactical Studies Rules Lake Geneva, WI.
- Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *To appear*.
- Lev Krasnov, Ivan Khokhlov, Maxim V Fedorov, and Sergey Sosnin. 2021. Transformer-based artificial neural networks for the conversion between chemical notations. *Scientific Reports*, 11(1):1–10.
- Pan Liu, Yanming Guo, Fenglei Wang, and Guohui Li. 2022. Chinese named entity recognition: The state of the art. *Neurocomputing*, 473:37–53.
- Behrang Mohit. 2014. Named entity recognition. *Natural language processing of semitic languages*, pages 221–245.
- Salman Naseer, Muhammad Mudasar Ghafoor, Sohaib bin Khalid Alvi, Anam Kiran, Shafique Ur Rahmand, Ghulam Murtazae, and Ghulam Murtaza. 2021. Named entity recognition (ner) in nlp techniques, tools accuracy and performance. *Pakistan Journal of Multidisciplinary Research*, 2(2):293–308.
- Thanh Vu Nguyen, Trung Nguyen, Vuong Son Ta, Pham Quang Nhat Minh, Vu Xuan Son, Nguyen Minh Hieu, Nguyen Phuong Thai, Vu Thanh Hung, Pham Minh Quang, Nguyen Thi Minh Huyen, and Haizhou Li. 2021. Trankit: A light-weight transformer-based toolkit for multilingual natural language processing. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*.
- Akila Peiris and Nisansa de Silva. 2022. Synthesis and evaluation of a domain-specific large data set for dungeons & dragons. *arXiv preprint arXiv:2212.09080*.
- Akila Peiris and Nisansa de Silva. 2023. SHADE: semantic hypernym annotator for Domain-Specific entities - DnD domain use case. In *2023 IEEE 17th International Conference on Industrial and Information Systems (ICIIS)*, page 6, Peradeniya, Sri Lanka.
- Christopher Perkins, Graeme Barber, Bill Benham, et al. 2021. *Candlekeep Mysteries*. Wizards of the Coast.
- Christopher Perkins, Will Doyle, and Steve Winter. 2017. *Tomb of Annihilation*. Wizards of the Coast.
- Christopher Perkins, Tracy Hickman, and Laura Hickman. 2016. *Curse of Strahd*. Wizards of the Coast.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Sungyong Seo, Sercan Arik, Jinsung Yoon, Xiang Zhang, Kihyuk Sohn, and Tomas Pfister. 2021. Controlling neural networks with rule representations. *Advances in Neural Information Processing Systems*, 34:11196–11207.

- Hemlata Shelar, Gagandeep Kaur, Neha Heda, and Poorva Agrawal. 2020. Named entity recognition approaches and their comparison for custom ner model. *Science & Technology Libraries*, 39(3):324–337.
- Nisansa de Silva and Dejing Dou. 2021. [Semantic oppositeness assisted deep contextual modeling for automatic rumor detection in social networks](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 405–415, Online. Association for Computational Linguistics.
- Nisansa de Silva, Dejing Dou, and Jingshan Huang. 2017. Discovering inconsistencies in pubmed abstracts through ontology-based information extraction. In *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 362–371.
- Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2019. Portuguese named entity recognition using bert-crf. *arXiv preprint arXiv:1909.10649*.
- Keet Sugathadasa, Buddhi Ayesha, Nisansa de Silva, Amal Shehan Perera, Vindula Jayawardana, Dimuthu Lakmal, and Madhavi Perera. 2017. Synergistic union of word2vec and lexicon for domain specific semantic similarity. In *2017 IEEE international conference on industrial and information systems (ICIIS)*, pages 1–6. IEEE.
- Simone Tedeschi, Valentino Maiorca, Niccolò Campolungo, Francesco Cecconi, and Roberto Navigli. 2021. WikiNEuRal: Combined neural and knowledge-based silver data creation for multilingual NER. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2521–2533.
- George R Terrell and David W Scott. 1992. Variable kernel density estimation. *The Annals of Statistics*, pages 1236–1265.
- 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.
- Peng Wang, Zhe Wang, Xiaowang Zhang, Kewen Wang, and Zhiyong Feng. 2021. Enhanced named entity recognition with semantic dependency. In *PRICAI 2021: Trends in Artificial Intelligence: 18th Pacific Rim International Conference on Artificial Intelligence, PRICAI 2021, Hanoi, Vietnam, November 8–12, 2021, Proceedings, Part II 18*, pages 287–298. Springer.
- Ralph Weischedel et al. 2013. Ontonotes release 5.0 ldc2013t19. Web Download.
- Yunzhi Yao, Shaohan Huang, Wenhui Wang, Li Dong, and Furu Wei. 2021. Adapt-and-distill: Developing small, fast and effective pretrained language models for domains. *arXiv preprint arXiv:2106.13474*.
- José P Zagal and Sebastian Deterding. 2018. Definitions of role-playing games. In *Role-Playing Game Studies*, pages 19–51. Routledge.
- Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi, and Prithviraj Ammanabrolu. 2022. An ai dungeon master’s guide: Learning to converse and guide with intents and theory-of-mind in dungeons and dragons. *arXiv preprint arXiv:2212.10060*.