

Unified Examination of Entity Linking in Absence of Candidate Sets

Nicolas Ong and Hassan S. Shavarani and Anoop Sarkar

School of Computing Science

Simon Fraser University

BC, Canada

{nong, sshavara, anoop}@sfu.ca

Abstract

Despite remarkable strides made in the development of entity linking systems in recent years, a comprehensive comparative analysis of these systems using a unified framework is notably absent. This paper addresses this oversight by introducing a new black-box benchmark and conducting a comprehensive evaluation of all state-of-the-art entity linking methods. We use an ablation study to investigate the impact of candidate sets on the performance of entity linking. Our findings uncover exactly how much such entity linking systems depend on candidate sets, and how much this limits the general applicability of each system. We present an alternative approach to candidate sets, demonstrating that leveraging the entire in-domain candidate set can serve as a viable substitute for certain models. We show the trade-off between less restrictive candidate sets, increased inference time and memory footprint for some models¹.

1 Introduction

Entity linking² is the task of annotating plain text with references to pre-defined entries - also known as entities - in a knowledge base. The annotation process involves finding specific parts of the text - also known as mentions - that may contain a knowledge base reference, and linking those to the right knowledge base record. For example, in the headline SOCCER - LATE GOALS GIVE JAPAN WIN OVER SYRIA, the entity linking system can link JAPAN (underlined) to *Japan national football team*.

Candidate generation is an integral part of many entity linking systems, in which a small set of candidate knowledge base entities are selected for each potential candidate span. Many recent entity

linking techniques assume the presence of the pre-existing candidate sets for all potential spans. The most popular such set is KB+Yago (Ganea and Hoffmann, 2017) which is used in a number of recent methods (Kolitsas et al., 2018; Peters et al., 2019; Poerner et al., 2020; Feng et al., 2022; van Hulst et al., 2020; De Cao et al., 2021b; Kannan Ravi et al., 2021; Shavarani and Sarkar, 2023). Another less popular set is PPRforNED (Perschina et al., 2015) which assumes the test sets are available to the annotators. Due to this assumption, very few recent publications evaluate using PPRforNED (Martins et al., 2019; De Cao et al., 2021b,a; Shavarani and Sarkar, 2023).

Unified evaluation of the different entity linking systems with respect to the application of candidate sets should play a crucial role in a better understanding of the strengths and weaknesses of each system. This will give the research community and commercial deployments better ways to select the most suitable system based on their needs while providing them a platform to identify avenues for enhancement.

In this paper, we propose a novel unified black-box evaluation framework for recent state-of-the-art entity linking systems. Our contributions in this paper are as follows:

(1) We unify the evaluation setup for the systems using GERBIL (Röder et al., 2018) and gerbil_connect (Shavarani and Sarkar, 2023), and black-box evaluate the systems over the same benchmark dataset CoNLL/AIDA (Hoffart et al., 2011) which allows us to abstract away their internal model structure and decoding algorithms.

(2) We discuss the importance of the pre-built candidate sets for obtaining good results on benchmarks in entity linking. However, candidate sets are not always available, and the literature lacks a systematic evaluation of the entity linking systems in absence of the candidate sets. To fill this gap, we suggest an experimental setup to replace them with

¹https://github.com/NicolasOng/gerbil_connects contains all the evaluated models as well as our evaluation source code.

²See Appendix A for an in-depth review of the modules in the entity linking pipeline.

a feasible set; the entire in-domain vocabulary of the benchmark dataset. Please note that our focus in these experiments is not to re-implement each technique, but rather to evaluate the *resilience* of the entity linking systems in absence of the carefully hand-crafted candidate sets.

(3) We examine the *adaptability* of the entity linking systems in the literature to unseen test data using the novel AIDA/testc dataset (Shavarani and Sarkar, 2023) which contains new annotations on news stories in 2020 with 924 novel entities.

2 Unified Black-Box Evaluation

We benchmark the recent entity linking systems, unchanged and as provided originally by their authors. In these experiments, we intend to examine the suitability of these systems as off-the-shelf systems which can be integrated in future downstream applications.

We unify the evaluation environment as GERBIL (Röder et al., 2018) plus `gerbil_connect` (Shavarani and Sarkar, 2023). In the evaluation procedure, GERBIL will provide the testing documents to `gerbil_connect` and receives the entity annotations in the format of (*begin char.*, *end char.*, *entity annotation*) from `gerbil_connect`. We implement `gerbil_connect` tailored to each entity linking system so that it can transform the evaluation documents to readable inputs for each system. Specifically, we have (1) utilized NLTK’s word tokenizer³ to transform raw non-tokenized evaluation sets into their expected CoNLL tokenized format for the models that depend on reading from AIDA test files (Peters et al., 2019; Poerner et al., 2020; Feng et al., 2022), (2) simulated long text splitting and result merging strategies for the models with input length constraints (Peters et al., 2019; Poerner et al., 2020; Feng et al., 2022; De Cao et al., 2021b), (3) implemented a subword token id to character id conversion for the models that output annotations as tokenized subword ids (Peters et al., 2019; Poerner et al., 2020; Feng et al., 2022), and (4) provided the external data sources such as the pre-built candidate sets to the model initializers where necessary (De Cao et al., 2021b,a). Appendix B provides more details on the mentioned alterations. Empirically, running the models without adding these techniques significantly hurts performance. Removing the tokenization step alone can drop the model performance by up to 20 Micro-F1 points.

³<https://www.nltk.org/api/nltk.tokenize.html>

Once done, `gerbil_connect` translates the produced annotations in each system back to the unified annotation format, understandable for GERBIL. We train the models that are not released by the authors (Poerner et al., 2020; Feng et al., 2022), using their own released source code, and do not consider the models which we were not able to acquire their training source code or were not able to get their training scripts to converge. (Martins et al., 2019; Févry et al., 2020; Mrini et al., 2022; Kannan Ravi et al., 2021; Broscheit, 2019; Xiao et al., 2023) Appendix C briefly discusses each of the evaluated systems. We direct the readers to the original publication of each model for a more in-depth understanding of their contributions.

We use CoNLL/AIDA evaluation sets `testa` and `testb` - reported by all entity linking systems tested in different evaluation frameworks - as well as the newly annotated AIDA/testc evaluation set. The results tables show the GERBIL InKB Micro-F1 evaluation results.

Table 1 presents the unified black-box evaluation results. The necessary unification adjustments mentioned above and the evaluation format has caused some evaluation scores to deviate from their original reported results. However, we have tried to control for this as much as possible. The Difference columns in Table 1 reflects on the mentioned score deviations.

In our experiments, we found that (Peters et al., 2019) suffered the most, with an approximate loss of 5% when comparing our results to the originally reported scores. The rest of the models were hit by at most 2%, confirming the reliability of our framework for further analysis. `testc` is a more challenging evaluation set which contains novel entities that typically hurt model recall. In our experimental results, we found that text generation models (from previous work) performed better on `testc`. However, the best performing model was not generative, but rather used structured prediction (Shavarani and Sarkar, 2023).

3 Candidate Set Ablations

Candidate sets are an integral part of entity linking systems, many of which assume the presence of good quality sets to perform well. Although this assumption holds when linking to English Wikipedia, it does not necessarily hold when considering other ontologies (e.g. UMLS; Bodenreider, 2004) and

	Micro-F1			Difference	
	testa	testb	testc	testa	testb
Kolitsas et al. (2018)	89.50	82.44	65.75	+0.10	+0.04
Peters et al. (2019)					
KnowBert-Wiki	76.74	71.68	54.12	-3.46	-2.72
KnowBert-W+W	77.19	71.69	53.92	-4.91	-2.01
Poerner et al. (2020)	89.40	84.83	65.93	-1.40	-0.17
van Hulst et al. (2020)					
Wiki 2014	83.30	82.53	71.69	-	-0.77
Wiki 2019	79.64	80.10	73.54	-	-0.40
De Cao et al. (2021b)	90.09	82.78	75.60	-	-0.92
De Cao et al. (2021a)	87.29	85.65	47.54	-	+0.15
Zhang et al. (2022)	86.81	84.30	72.55	-	-1.50
Feng et al. (2022)	87.64	86.49	65.05	-	+0.19
Shavarani and Sarkar (2023)					
large-500K (no cnds.)	89.72	82.25	77.54	+0.02	+0.05
large-500K (Kb+Yago)	89.89	82.88	59.50	+0.09	+0.08
large-500K (PPRforNED)	91.58	85.22	46.98	+0.08	+0.02

Table 1: Comparison of recent entity linking systems within the unified black-box testing framework of GERBIL + gerbil_connect. Difference column reports the difference between our unified evaluation environment and the originally reported numbers. We have assessed all models twice for consistency. Except for (De Cao et al., 2021b), all models yielded identical scores, while De Cao et al. (2021b) showed a low variance of 0.08 in the results. Thus, the results mirror those reported by GERBIL, with the exception of (De Cao et al., 2021b), which is averaged over two runs.

languages⁴.

In this section, we ablate the mention-specific candidate sets from the entity linking systems to study their performance in absence of the hand-crafted candidate sets. For our experiments, we selected the candidate-set-independent setting of the models in any system that provides such a setting. For the other systems that require a candidate set, and we cannot remove the candidate set dependence, we return the entire in-domain mention vocabulary of AIDA (the in-domain fixed candidate set of Shavarani and Sarkar, 2023) as the replacement for the required candidate sets (5598 entities including the *None* entity). Where applicable, we add priors such that each candidate has an equal probability.

Table 2 demonstrates the evaluation results of the models with Micro-F1 scores above 1.0 after considering the candidate-independent version of the models, or the candidate set expansion.

We experimented with removing candidate sets altogether, but the models that appear in Table 1 but do not appear in Table 2 essentially failed without candidate sets, resulting in Micro-F1 scores

		Micro-F1		
		testa	testb	testc
a)	De Cao et al. (2021b)	85.15	78.98	75.62
	De Cao et al. (2021a)	62.00	49.51	37.05
	Zhang et al. (2022)	86.81	84.30	72.55
	Shavarani and Sarkar (2023)	89.72	82.25	77.54
b)	Poerner et al. (2020)	22.81	18.81	17.56
	Feng et al. (2022)	35.00	32.58	27.48

Table 2: Comparison of entity linking systems after a) running the model with no access to hand-crafted candidate sets b) modifying the model to consider the entire AIDA in-domain vocabulary as the candidate set.

close to 0. These results demonstrate that most entity linking systems are too intertwined with their candidate sets and without this additional data resource, the systems do not produce useful results and are too brittle to be used in real-world production deployments.

Table 2 results prove that generation-based systems are more *resilient* against candidate sets. Nonetheless, without given candidate sets, De Cao et al. (2021b) and De Cao et al. (2021a) lose approximately 5% and 20-30% of their best performance, respectively. Shavarani and Sarkar (2023) -

⁴See Botha et al. (2020) for more discussion.




Figure 1: Entity linking error distribution in four categories of over-generated (gray, vertical), under-generated (red, horizontal), incorrect entity (teal, north east) and incorrect mention (blue, north west) before candidate set ablations (left) and after the ablations (right). The y-axis is the error analysis ratio as described below.

a non-generation-based system, designed without dependence on candidate sets and only using these resources to improve performance - suffers the least and loses only 2% of its best performance without candidate sets.

The largest performance drop in our experiments correlates with using mention-entity similarity methods for entity disambiguation, where a representation of the mention and entity are compared to determine the most relevant entity.

In these systems, models that generate mention representations by combining candidate entity representations see their performance decreased to 20%-35%, while models that generate mention representations by combining the word or token representations within or surrounding the mention perform too poorly to be present in Table 2.

Shavarani and Sarkar (2023) and De Cao et al. (2021b) only show an approximate 2% drop in performance, showing that they can easily handle a larger set of candidate entities.

The larger candidate sets lead to longer inference times. The run time for Feng et al. (2022); Kolitsas et al. (2018); Poerner et al. (2020); Peters et al. (2019) that compare the mentions to each entity in the candidate set increases by 90x, 50x, 25x, and 10x, respectively. van Hulst et al. (2020) does not follow this trend since it selects the 30 candidate entities with the highest prior before performing entity disambiguation.

Error Analysis. We store the produced annotations from each system reported in Table 1 (w/ candidate sets) and Table 2 (w/o candidate sets), and compare their produced annotations with the expected annotations of AIDA/testa (4791 annotations). For models with multiple reported settings, we select the setting correlated to AIDA/testc as it represents the most generalization-capable setting for unseen in-domain documents.

We count the number of annotations in four error categories of over-generated, under-generated, incorrect mention and incorrect entity, and divide each by the total number of gold annotations. Figure 1 presents the calculated error analysis ratios. Over-generation refers to annotations made by the model and not in the gold set. Under-generation refers to annotations in the gold set but not made by the model. Incorrect entity refers to annotations where the model linked the wrong entity. Incorrect mention refers to annotations where the span’s start or end is incorrect.

Before ablation of candidate sets (Figure 1-left), (van Hulst et al., 2020) has the highest rate of over-generation followed by (Zhang et al., 2022), while (Peters et al., 2019) shows the lowest over-generation rate. On the other hand (Peters et al., 2019) has the highest under-generation ratio establishing itself as a conservative entity linking system.

Comparing the performance of entity systems w/

and w/o candidate sets, the biggest increase is seen in incorrect entity prediction ratios, confirming the dependence of entity linking systems to candidate sets. (Poerner et al., 2020) sees the biggest increase in incorrect entity predictions followed by (Feng et al., 2022). While (Zhang et al., 2022) and (Shavarani and Sarkar, 2023) report the smallest rate increase in this category as these methods are less dependent on candidate sets. (Feng et al., 2022) on the other hand shows an increase in under-generation signaling the effect of candidate sets in prediction confidence for this system.




Figure 2: Entity linking micro precision (blue, north east) and recall (red, north west) score differences over tests between model’s original configuration and candidate set ablation configuration.

Figure 2 illustrates the disparities in precision and recall pre- and post-ablation of candidate sets for the models outlined in Table 2. Our findings reveal that candidate sets significantly enhance precision and recall. With the exception of (Zhang et al., 2022), which generates candidates in real-time, the other systems show that without candidate sets there is a substantial decrease in precision and recall, exceeding 60% for (Poerner et al., 2020).

4 Conclusion

In this study, we have successfully established a unified black-box evaluation framework for modern entity linking techniques. We provide an in-depth ablation analysis to examine the significance of hand-crafted candidate sets for the entity linking task. Our findings confirm that modern entity linking systems are excessively dependent on candidate sets. Our study shows that we need to be less reliant on hand-crafted candidate sets in order to ensure robust, versatile and accurate entity linking systems.

References

- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. [Contextual string embeddings for sequence labeling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Hannah Bast, Matthias Hertel, and Natalie Prange. 2022. [ELEVANT: A fully automatic fine-grained entity linking evaluation and analysis tool](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 72–79, Abu Dhabi, UAE. Association for Computational Linguistics.
- Olivier Bodenreider. 2004. [The unified medical language system \(umls\): integrating biomedical terminology](#). *Nucleic acids research*, 32(suppl_1):D267–D270.
- Jan A. Botha, Zifei Shan, and Daniel Gillick. 2020. [Entity Linking in 100 Languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7833–7845, Online. Association for Computational Linguistics.
- Samuel Broscheit. 2019. [Investigating entity knowledge in BERT with simple neural end-to-end entity linking](#). pages 677–685.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021a. [Highly parallel autoregressive entity linking with discriminative correction](#). pages 7662–7669.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021b. [Autoregressive entity retrieval](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yukun Feng, Amir Fayazi, Abhinav Rastogi, and Manabu Okumura. 2022. [Efficient entity embedding construction from type knowledge for BERT](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 1–10, Online only. Association for Computational Linguistics.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. [Incorporating non-local information into information extraction systems by Gibbs sampling](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 363–370, Ann Arbor, Michigan. Association for Computational Linguistics.
- Thibault Févry, Nicholas Arthur FitzGerald, Livio Baldini Soares, and Tom Kwiatkowski. 2020. [Empirical](#)

- evaluation of pretraining strategies for supervised entity linking. In *AKBC 2020 - Automated Knowledge Base Construction*.
- Octavian-Eugen Ganea and Thomas Hofmann. 2017. **Deep joint entity disambiguation with local neural attention**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2619–2629, Copenhagen, Denmark. Association for Computational Linguistics.
- 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.
- Manoj Prabhakar Kannan Ravi, Kuldeep Singh, Isaiah Onando Mulang’, Saeedeh Shekarpour, Johannes Hoffart, and Jens Lehmann. 2021. **CHOLAN: A modular approach for neural entity linking on Wikipedia and Wikidata**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 504–514, Online. Association for Computational Linguistics.
- 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.
- Pedro Henrique Martins, Zita Marinho, and André F. T. Martins. 2019. **Joint learning of named entity recognition and entity linking**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 190–196, Florence, Italy. Association for Computational Linguistics.
- Khalil Mrini, Shaoliang Nie, Jiatao Gu, Sinong Wang, Maziar Sanjabi, and Hamed Firooz. 2022. **Detection, disambiguation, re-ranking: Autoregressive entity linking as a multi-task problem**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1972–1983, Dublin, Ireland. Association for Computational Linguistics.
- Maria Pershina, Yifan He, and Ralph Grishman. 2015. **Personalized page rank for named entity disambiguation**. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 238–243, Denver, Colorado. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. 2019. **Knowledge enhanced contextual word representations**. pages 43–54.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2020. **E-BERT: Efficient-yet-effective entity embeddings for BERT**. pages 803–818.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- Michael Röder, Ricardo Usbeck, and Axel-Cyrille Ngonga Ngomo. 2018. **GERBIL - benchmarking named entity recognition and linking consistently**. *Semantic Web*, 9(5):605–625.
- Ahmad Sakor, Isaiah Onando Mulang’, Kuldeep Singh, Saeedeh Shekarpour, Maria Esther Vidal, Jens Lehmann, and Sören Auer. 2019. **Old is gold: Linguistic driven approach for entity and relation linking of short text**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2336–2346, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hassan Shavarani and Anoop Sarkar. 2023. **SpEL: Structured prediction for entity linking**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11123–11137, Singapore. Association for Computational Linguistics.
- Johannes M van Hulst, Faegheh Hasibi, Koen Dercksen, Krisztian Balog, and Arjen P de Vries. 2020. **Rel: An entity linker standing on the shoulders of giants**. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2197–2200.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. **Scalable zero-shot entity linking with dense entity retrieval**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6397–6407, Online. Association for Computational Linguistics.
- Zilin Xiao, Ming Gong, Jie Wu, Xingyao Zhang, Linjun Shou, Jian Pei, and Daxin Jiang. 2023. **Instructed language models with retrievers are powerful entity linkers**.
- Ikuya Yamada, Akari Asai, Jin Sakuma, Hiroyuki Shindo, Hideaki Takeda, Yoshiyasu Takefuji, and Yuji Matsumoto. 2020. **Wikipedia2Vec: An efficient toolkit for learning and visualizing the embeddings of words and entities from Wikipedia**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 23–30. Association for Computational Linguistics.
- Wenzheng Zhang, Wenyue Hua, and Karl Stratos. 2022. **Entqa: Entity linking as question answering**. *International Conference on Learning Representations*.

A A More In-Depth Literature Review

Entity linking can be viewed as a multi-task pipeline which performs mention detection, candidate generation, and entity disambiguation (see Broscheit, 2019; van Hulst et al., 2020; Shavarani and Sarkar, 2023). *Mention detection* focuses on the discovery of potential references in the given text to knowledge base entities. *Candidate Generation* short-lists a small set of candidate knowledge base entities for each potential span. *Entity Disambiguation* selects one of the short-listed entities to link to the selected span.

Mention Detection. Common approaches include: (1) listing every possible span of length up to n (usually 5) tokens (Kolitsas et al., 2018; Peters et al., 2019; Poerner et al., 2020; Feng et al., 2022), (2) using a named entity recognition (NER) tool (van Hulst et al., 2020; Hoffart et al., 2011) such as Flair (Akbik et al., 2018) or the Stanford NER Tagger (Finkel et al., 2005), (3) scoring each token as a potential mention and merging the consecutive predictions referring to the same entity (Broscheit, 2019; Shavarani and Sarkar, 2023), and (4) probabilistic prediction of tokens beginning, ending, or outside (BIO) of mentions (Kannan Ravi et al., 2021; Févry et al., 2020; De Cao et al., 2021a; Zhang et al., 2022; Xiao et al., 2023).

Candidate Generation. This module selects the most relevant of all possible entities for a candidate span in the input text. As previously mentioned, most models rely on pre-existing candidate sets such as KB+Yago (Ganea and Hofmann, 2017). Kannan Ravi et al. (2021) use the BM25 (Robertson et al., 1995) algorithm to rank entities for a mention. Zhang et al. (2022) and Xiao et al. (2023) explore document-level candidate sets that aid the mention detection step. Both rank the dot product scores of entity embeddings by an input document embedding. Zhang et al. (2022) use a fine-tuned BLINK (Wu et al., 2020) model. In the extreme case, some methods assume the entire entity vocabulary as the candidate set (Broscheit, 2019; Févry et al., 2020; Shavarani and Sarkar, 2023). Such methods usually limit the size of the entity vocabulary to 500K-700K of the most frequent entities. Wikipedia contains over 6.5M entities. In candidate generation, mentions with an empty candidate set or the ones with the most probable predicted entities being either the *None* entity or not in the candidate set, will be automatically ignored. Thus candidate sets help prevent over-generation and

improve the overall quality of entity linking.

Entity Disambiguation. Common methods include: (1) employing constrained beam search to generate the entities with the candidate sets constraining the search process (De Cao et al., 2021b; Xiao et al., 2023), (2) *argmax* selection over the predicted probability distributions over the candidate set (Broscheit, 2019; De Cao et al., 2021a; Shavarani and Sarkar, 2023), and (3) ranking the entity candidates using a similarity metric which scores (entity prediction, mention) representation pairs.

For the latter method, the mention representations can be generated by combining the character, word, or token embeddings within the mention span (Kolitsas et al., 2018; Peters et al., 2019; Martins et al., 2019; Févry et al., 2020), or of the context surrounding the mention (van Hulst et al., 2020; Martins et al., 2019), or all the embeddings of the candidates for the mention (Poerner et al., 2020; Feng et al., 2022). The entity representations can come from a pre-computed source like KB+Yago (Ganea and Hofmann, 2017) in (Kolitsas et al., 2018; van Hulst et al., 2020; Peters et al., 2019), Wikipedia2Vec (Yamada et al., 2020) in (van Hulst et al., 2020; Poerner et al., 2020), or PPRforNED (Perschina et al., 2015) in (Martins et al., 2019). As well, the entity representations can be textual documents with items like the entity’s title, aliases, and description (Kannan Ravi et al., 2021; Hoffart et al., 2011) or model representations of such documents (Feng et al., 2022; Févry et al., 2020). Dot product is most commonly used to score mention-entity similarity (Kolitsas et al., 2018; Peters et al., 2019; Poerner et al., 2020; Feng et al., 2022; Févry et al., 2020). For representations in the form of natural language (as opposed to dense embedding vectors), word overlap, KL-divergence, and n-gram based measures (Hoffart et al., 2011), as well as document similarity prediction using fine-tuned PLMs (Kannan Ravi et al., 2021) have been used.

Methods using mention-entity similarity tend to also use mention-entity priors (i.e., $\hat{p}(e|m)$) (Kolitsas et al., 2018; Peters et al., 2019; van Hulst et al., 2020; Poerner et al., 2020; Feng et al., 2022) and thematic coherence (prioritizing entities similar to other entities in the document) (Hoffart et al., 2011; Kolitsas et al., 2018; van Hulst et al., 2020) to bias the results. The priors are commonly taken from Ganea and Hofmann (2017).

End-to-end entity linking systems usually combine the mention detection and entity disambigua-

tion steps through sharing the same underlying pre-trained language model encoders (Kolitsas et al., 2018; Broscheit, 2019; Shavarani and Sarkar, 2023). Xiao et al. (2023) also combines these steps with their in-context learning approach, which queries a LLM (such as GPT-3, GPT-3.5, or Codex) with a natural language prompt containing the task, candidate entities, and the text to be linked.

B Model Changes for Unified Black Box Evaluation

As mentioned, we use GERBIL to evaluate the models, and `gerbil_connect` acts as the middle-ware between GERBIL and the models. GERBIL provides the model with raw un-tokenized text and expects annotations in the format (*begin char., end char., entity annotation*). In the unification process leading to the specific implementations of `gerbil_connect`, we altered each reproduced model to enable them to work within our unified black-box evaluation environment. The following sections discuss each alternation in more details. Table 3 discusses the required alterations to adapt each model to our evaluation framework.

We encountered limitations in assessing certain models within our framework. The models from Martins et al. (2019); Févry et al. (2020); Mrini et al. (2022) were not included due to the unavailability of their source code. The model from Kannan Ravi et al. (2021) was also excluded, as we lacked access to their pre-trained model and the training script was not provided. The models from Broscheit (2019); Xiao et al. (2023) were not considered due to the absence of their pre-trained models and the extended duration of their training process, which was not feasible for inclusion in the manuscript.

For the purpose of facilitating replication and standardizing the assessment of emerging entity linking methodologies in subsequent studies, we strongly advise authors to make their source code, trained models, and a `gerbil_connect` integration with GERBIL publicly available. In the absence of these, we suggest the provision of a straightforward function that accepts raw text and outputs a list of annotations. This approach would streamline the process of incorporating an entity linking technique into any evaluation setting.

	Modifications
Kolitsas et al. (2018)	-
Peters et al. (2019)	1, 2, 3
Poerner et al. (2020)	1, 2, 3, 5
van Hulst et al. (2020)	-
De Cao et al. (2021b)	2, 4
De Cao et al. (2021a)	4
Zhang et al. (2022)	-
Feng et al. (2022)	1, 2, 3, 5
Shavarani and Sarkar (2023)	-

Table 3: Comparison of major modifications made to each system to fit into the unified evaluation environment.

B.1 Input Tokenization

Some models require their input to be tokenized, but do not have this implemented. Peters et al. (2019) reads directly from the AIDA dataset file where each token is on a new line. Our evaluation environment gives the model text, so we add a tokenization step before the model so the text is closer to the expected input. We used NLTK’s word tokenizer for this, as it tokenizes the text similar to the expected CoNLL tokenized format. After the model makes annotations for each of these word-level tokens, we add another step at the end to convert the predictions from word level to character level.

B.2 Document Splitting

Many models have a length limit for the documents they annotate, and therefore need a document splitting strategy. De Cao et al. (2021b) uses a document splitting strategy during evaluation⁵, however this strategy is not described in the paper and the code is not in their released repository. We used the splitting strategy created by Bast et al. (2022) who were able to replicate De Cao et al. (2021b).

We also add a step at the end to map the annotations for each section back into the original document.

B.3 Token-to-Character Annotations

Many models convert the gold annotations to token-level, and evaluate on their token-level predictions. This allows them to forgo converting their token-level annotations back into character-level annotations. However, our evaluation environment requires character-level annotations as output, so we

⁵see <https://github.com/facebookresearch/GENRE/issues/30>

add a step to perform this conversion where necessary.

B.4 Outside or Custom Data

The data used by most models to perform entity linking are available to download. However, De Cao et al. (2021b) requires custom data that hasn't been released⁶. To fix this, we used the data created by Bast et al. (2022). De Cao et al. (2021a) also requires custom data created by the authors, but this data has been provided in their released repository.

B.5 Training a New Model

Although some publications have not released their pre-trained checkpoints, they have released their training data and scripts. In such cases, we trained new model checkpoints to reproduce the original results.

C Evaluated Model Descriptions

Here we provide overviews of the evaluated models in our benchmarking experiments. For brevity, we don't discuss every detail here, and solely rely on the important design choices in each method.

C.1 Traditional Models

Traditional models rely on pre-computed dictionaries such as KB+Yago (Ganea and Hofmann, 2017) to provide candidate sets for detected mentions. Replacing the dictionary with a static candidate set for all mentions - such as an empty set or a large set with 5K entities results in poor performance.

C.1.1 Kolitsas et al. (2018)

For mention detection and candidate generation, Kolitsas et al. (2018) use KB+Yago (Ganea and Hofmann, 2017) on all n-length spans in the input text. For entity disambiguation, Kolitsas et al. (2018) generate mention representations from the tokens in each mention, then use dot-product similarity to rank each candidate entity using Ganea and Hofmann (2017) entity representations. In addition, mention-entity priors and thematic coherence are used to bias each candidate entity's score.

C.1.2 Peters et al. (2019)

For mention detection and candidate generation, Peters et al. (2019) use KB+Yago (Ganea and Hofmann, 2017) on all n-length spans in the input text.

⁶see <https://github.com/facebookresearch/GENRE/issues/37> and Appendix A.2 Setting in De Cao et al. (2021b).

For entity disambiguation, Peters et al. (2019) generate mention representations from the word pieces in each span, then use dot-product similarity to rank each candidate entity using Ganea and Hofmann (2017) entity representations. In addition, mention-entity priors are used to bias each candidate entity's score.

C.1.3 van Hulst et al. (2020)

For mention detection, van Hulst et al. (2020) use Flair (Akbik et al., 2018) NER annotation tool. Mention representations are generated using their surrounding 50-word context. For candidate generation, van Hulst et al. (2020) use KB+Yago (Ganea and Hofmann, 2017). Then, they refine the candidate list by selecting the 4 entities with the highest priors and 3 with the highest dot-product similarity between the mention representation and their Ganea and Hofmann (2017) embedding. For entity disambiguation, van Hulst et al. (2020) attempt to maximize the sum of *context coherence* (a coherence score between an entity and the mention's local context) and thematic coherence for all mentions. The following providing the context coherence calculation formula:

$$\psi(e, c) = \sum_{w \in c} \beta(w) x_w^\top B x_w$$

where c is context (made up of words w), e is the entity, $\beta(w)$ is a weight for each word and B is a trainable matrix.

C.1.4 Poerner et al. (2020)

For mention detection and candidate generation, Poerner et al. (2020) use KB+Yago (Ganea and Hofmann, 2017) on all n-length spans in the input text. For entity disambiguation, entity candidates are ranked using dot-product similarity. Entity representations are retrieved from Wikipedia2Vec (Yamada et al., 2020). Mention representations are generated using a BERT model (Devlin et al., 2019). The mention span in the input text is masked by the sum of its candidate entity representations. The mention representation is the mask's embedding generated by the BERT model. In addition, mention-entity priors are used to bias the results.

C.1.5 Kannan Ravi et al. (2021)

For mention detection, each token is classified using a BERT model fine-tuned for BIO tagging. For candidate generation, Kannan Ravi et al. (2021) use KB+Yago (Ganea and Hofmann, 2017) and the BM25 algorithm (Robertson et al., 1995) to query a

local index (Sakor et al., 2019) of items from Wikidata. For entity disambiguation, Kannan Ravi et al. (2021) use a BERT model *WikiBERT* fine-tuned to predict document similarity. Wikipedia articles are used as the entity documents, and the local context of the mention in the text is used as the mention document.

C.1.6 Feng et al. (2022)

For mention detection and candidate generation, Feng et al. (2022) use KB+Yago (Ganea and Hofmann, 2017) on all n-length spans in the input text. For entity disambiguation, entity candidates are ranked using dot-product similarity. Entity representations are generated by encoding the labels of other entities connected through *instance_of* edges in the KB. Mention representations are generated using a BERT model (Devlin et al., 2019). The mention span in the input text is masked by the sum of its candidate entity representations. The mention representation is the mask’s embedding generated by the BERT model. In addition, mention-entity priors are used to bias the results.

C.2 Generative Models

Generative models are able to generate candidate sets without relying on a pre-computed mention-to-candidates dictionary. This makes them more robust in settings where these mention-specific candidate sets are lacking. However, some can still use these candidate sets to improve results.

C.2.1 De Cao et al. (2021b)

De Cao et al. (2021b) output an annotated text of the input by generating it token-by-token. At every token, the model decides either to continue the input or to start a mention. Once inside a mention, the model decides at what token the mention span ends. KB+Yago (Ganea and Hofmann, 2017) is used to obtain a modified candidate set. For entity disambiguation, a candidate is generated using beam search constrained with a trie generated from the candidate entities.

C.2.2 De Cao et al. (2021a)

For mention detection, De Cao et al. (2021a) classify each token as the start or end of a mention. For candidate generation, De Cao et al. (2021a) use a given candidate set for each mention. For entity disambiguation, the model scores each candidate with an LSTM then uses a classifier to re-rank the candidates. When not provided with the candidate

sets, the model can use the LSTM to generate candidates using constrained beam search.

C.2.3 Zhang et al. (2022)

For candidate generation, Zhang et al. (2022) generate document-level candidates by selecting the top-K entities based on dot-product similarity using an encoder to create an encoding for both the input document and all the entities in Wikipedia (approximately 6M). For mention detection, it classifies each token as a start or end of a mention for each entity. For entity disambiguation, it chooses the entity with the start and end probability scores above a certain threshold.

C.2.4 Xiao et al. (2023)

For candidate generation, Xiao et al. (2023) use dot product similarity to rank entity encodings against the input document encoding, both generated using fine-tuned BERT (Devlin et al., 2019) text encoders. For mention detection, it generates a list of possible surface forms for each candidate entity using a modified KB+Yago (Ganea and Hofmann, 2017) dictionary. Then it searches the text for the generated surface forms. For entity disambiguation, it uses constrained beam search to both decide if each candidate mention is valid, and which entity is the correct entity for that mention.

They also experiment with another approach, *INSGENEL-ICL*, which sends a natural language prompt to a LLM with the task description, input document, and candidate entities, and expects a list of annotations as an answer.

C.3 Structured Prediction Based Models

These models don’t rely on mention-specific candidate sets or need to generate their own. They disambiguate among the entire entity vocabulary. However, some can still use candidate sets to improve results.

C.3.1 Broscheit (2019)

Broscheit (2019) performs entity linking as a per-token classification task over the entire entity vocabulary (the 700K most frequent entities in English Wikipedia). They fine-tune a BERT model for this task.

C.3.2 Shavarani and Sarkar (2023)

Shavarani and Sarkar (2023) model entity linking as structured prediction with a variable classification vocabulary size. The top-k entities for each

token are collected, and merged with adjacent tokens belonging to the same span. When a mention-specific candidate set is available, entities not in this candidate set are filtered out from the mention's candidate entities.