ENTYFI: A System for Fine-grained Entity Typing in Fictional Texts

Cuong Xuan Chu **Simon Razniewski** **Gerhard Weikum**
Max Planck Institute for Informatics
Saarbrücken, Germany
{cxchu, srazniew, weikum}@mpi-inf.mpg.de

Abstract

Fiction and fantasy are archetypes of long-tail domains that lack suitable NLP methodologies and tools. We present ENTYFI, a web-based system for fine-grained typing of entity mentions in fictional texts. It builds on 205 automatically induced high-quality type systems for popular fictional domains, and provides recommendations towards reference type systems for given input texts. Users can exploit the richness and diversity of these reference type systems for fine-grained supervised typing, in addition, they can choose among and combine four other typing modules: pre-trained real-world models, unsupervised dependency-based typing, knowledge base lookups, and constraint-based candidate consolidation. The demonstrator is available at https://d5demos.mpi-inf.mpg.de/entyfi.

1 Introduction

Motivation and Problem. Entity types are a core building block of current knowledge bases (KBs) and valuable for many natural language processing tasks, such as coreference resolution, relation extraction and question answering (Lee et al., 2006; Carlson et al., 2010; Recasens et al., 2013). Context-based entity typing, the task of assigning semantic types for mentions of entities in textual contexts (e.g., `musician`, `politician`, `location` or `battle`) therefore has become an important NLP task. While traditional methods often use coarse-grained classes, such as `person`, `location`, `organization` and `misc`, as targets, recent methods try to classify entities into finer-grained types, from hundreds to thousands of them, yet all limited to variants of the real world, like from Wikipedia or news (Lee et al., 2006; Ling and Weld, 2012; Corro et al., 2015; Choi et al., 2018).

Entity type information plays an even more important role in literary texts from fictional domains. Fiction and fantasy are core parts of human culture, spanning from traditional folks and myths into books, movies, TV series and games. People have created sophisticated fictional universes such as the Marvel Universe, DC Comics, Middle Earth or Harry Potter. These universes include entities, social structures, and events that are completely different from the real world. Appropriate entity typing for these universes is a prerequisite for several end-user applications. For example, a Game of Thrones fan may want to query for House Stark members who are Faceless Men or which character is both a Warg and a Greenseer. On the other hand, an analyst may want to compare social structures between different mythologies or formations of different civilizations.

State-of-the-art methods for entity typing mostly use supervised models trained on Wikipedia content, and only focus on news and similar real-world texts. Due to low coverage of Wikipedia on fictional domains, these methods are thus not sufficient for literary texts. For example, for the following sentence from *Lord of the Rings:*

> *"After Melkor's defeat in the First Age, Sauron became the second Dark Lord and strove to conquer Arda by creating the Rings"*

state-of-the-art entity typing methods only return few coarse types for entities, such as `person` for SAURON and MELKOR or `location` for FIRST AGE and ARDA. Moreover, existing methods typically produce predictions for each individual mention, so that different mentions of the same entity may be assigned incompatible types, e.g., ARDA may be predicted as `person` and `location` in different contexts.

Contribution. The prototype system presented in this demo paper, ENTYFI (fine-grained ENtity TYping on FIctional texts, see Chu et al. (2020) for full details) overcomes the outlined limitations.

Proceedings of the 2020 EMNLP (Systems Demonstrations), pages 100–106
November 16-20, 2020. ©2020 Association for Computational Linguistics

ENTYFI supports long input texts from any kind of literature, as well as texts from standard domains (e.g., news). With the sample text above, ENTYFI is able to predict more specific and meaningful types for entity mentions:

MELKOR:	`Ainur, Villain`	FIRST AGE:	`Eras, Time`
SAURON:	`Maiar, Villain`	DARK LORD:	`Ainur, Title`
RINGS:	`Jewelry, Magic Things`	ARDA:	`Location`

To address the lack of reference types, ENTYFI leverages the content of fan communities on Wikia.com, from which 205 reference type systems are induced. Given an input text, ENTYFI then retrieves the most relevant reference type systems and uses them as typing targets. By combining supervised typing method with unsupervised pattern extraction and knowledge base lookups, suitable type candidates are identified. To resolve inconsistencies among candidates, ENTYFI utilizes an integer linear programming (ILP) based consolidation stage.

Outline. The following section describes the architecture of ENTYFI with the approach underlying its main components. The demonstration is illustrated afterwards through its graphical user interface. Our demonstration system is available at: `https://d5demos.mpi-inf.mpg.de/entyfi`. We also provide a screencast video demonstrating our system, at: `https://youtu.be/g_ESaONagFQ`.

2 System Overview

ENTYFI comprises five steps: type system construction, reference universe ranking, mention detection, mention typing and type consolidation. Figure 1 shows an overview of the ENTYFI architecture.

2.1 Type System Construction

To counter the low coverage of entities and relevant types in Wikipedia for fictional domains, we make use of an alternative semi-structured resource, Wikia[1].

Wikia. Wikia is a large fiction community portal, includes over 385,000 individual universes. It covers a wide range of universes in fiction and fantasy, from old folks and myths like *Greek Mythology, Egyptian Mythology* to recent stories like *Harry Potter, The Lord of the Rings*. It also contains popular movies, TV series (e.g. *Game of Thrones, Breaking Bad*) and video games (e.g. *League of Legends, Pokemon*).

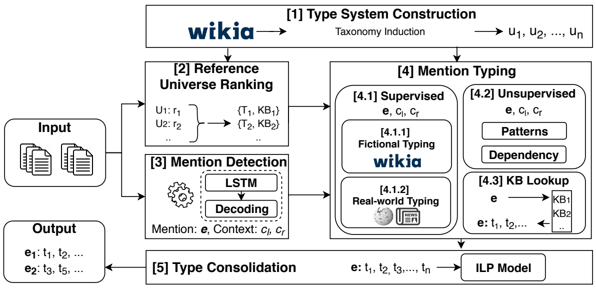

Figure 1: Overview of the architecture of ENTYFI (Chu et al., 2020).

Each universe in Wikia is organized similarly to Wikipedia, such that they contain entities and categories that can be used to distill reference type systems. We adopt techniques from the TiFi system (Chu et al., 2019) to clean and structure Wikia categories. We remove noisy categories (e.g. meta-categories) by using a dictionary-based method. To ensure connectedness of taxonomies, we integrate the category networks with WordNet (WN) by linking the categories to the most similar WN synsets. The similarity is computed between the context of the category (e.g., description, super/sub categories) and the gloss of the WN synset (Chu et al., 2019). Resulting type systems typically contain between 700 to 10,000 types per universe.

2.2 Reference Universe Ranking

Given an input text, the goal of this step is to find the most relevant universes among the reference universes. Each reference universe is represented by its entities and entity type system. We compute the cosine similarity between the TF-IDF vectors of the input and each universe. The top-ranked reference universes and their type systems are then used for mention typing (section 2.4).

2.3 Mention Detection

To detect entity mentions in the input text, we rely on a BIOES tagging scheme. Inspired by He et al. (2017) from the field of semantic role labeling, we design a BiLSTM network with embeddings and POS tags as input, highway connections between layers to avoid vanishing gradients (Zhang et al., 2016), and recurrent dropout to avoid over-fitting (Gal and Ghahramani, 2016). The output is then put into a decoding step by using dynamic programming to select the tag sequence with maximum score that satisfies the BIOES constraints. The de-

[1] https://wikia.com

coding step does not add more complexity to the training.

2.4 Mention Typing

We produce type candidates for mentions by using a combination of supervised, unsupervised and lookup approaches.

Supervised Fiction Types. Given an entity mention and its textual context, we approach typing as multiclass classification problem. The mention representation is the average of all embeddings of tokens in the mention. The context representation is a combination of left and right context around the mention. The contexts are encoded by using BiLSTM models (Graves, 2012) and then put into attention layer to learn the weight factors (Shimaoka et al., 2017). Mention and context representations are concatenated and passed to the final logistic regression layer with cross entropy loss function to predict the type candidates.

Target Classes. There are two kinds of target classes: (i) general types - 7 disjunct high-level WordNet types that we manually chose, mirroring existing coarse typing systems: `living_thing`, `location`, `organization`, `object`, `time`, `event`, `substance`, (ii) top-performing types - types from reference type systems. Due to a large number of types as well as insufficient training data, predicting all types in the type systems is not effective. Therefore, for each reference universe, we predict those types for which, on withheld test data, at least 0.8 F1-score was achieved. This results in an average of 75 types per reference universe.

Supervised Real-world Types. Although fictional universes contain fantasy contents, many of them reflect our real-world, for instance, *House of Cards*, a satire of American politics. Even fictional stories like *Game of Thrones* or *Lord of the Rings* contain types presented in real world, such as `King` or `Battle`. To leverage this overlap, we incorporate the Wikipedia- and news-trained typing model from Choi et al. (2018), which is able to predict up to 10,331 real-world types.

Unsupervised Typing. Along with supervised technique, we use a pattern-based method to extract type candidates which appear explicitly in contexts for mentions. We use 36 manually crafted Hearst-style patterns for type extraction (Seitner et al., 2016). Moreover, from dependency parsing, a noun phrase can be considered as a type candidate if there exists a *noun compound modifier (nn)* between the noun phrase and the given mention. In the case of candidate types appearing in the mention itself, we extract the head word of the mention and consider it as a candidate if it appears as a noun in WordNet. For example, given the text *Queen Cersei was the twentieth ruler of the Seven Kingdoms*, `queen` and `kingdom` are type candidates for the mentions CERSEI and SEVEN KINGDOMS, respectively.

KB Lookup. Using top-ranked universes from section 2.2 as basis for the lookup, we map entity mentions to entities in reference universes by using lexical matching. The types of entities in corresponding type systems then become type candidates for the given mentions.

2.5 Type Consolidation

Using multiple universes as reference and typing in long texts may produce incompatibilities in predictions. For example, SARUMAN, a wizard in Lord of the Rings can be predicted as a white walker using the model learnt from Game of Thrones. To resolve possible inconsistencies, we rely on a consolidation step that uses an integer linear programming (ILP) model. The model captures several constraints, including disjointness, hierarchical coherence, cardinality limit and soft correlations (Chu et al., 2020).

ILP Model. Given an entity mention ***e*** with a list of type candidates with corresponding weights, a decision variable T_i is defined for each type candidate t_i. $T_i = 1$ if ***e*** belongs to t_i, otherwise, $T_i = 0$. With the constraints mentioned above, the objective function is:

maximize

$$\alpha \sum_i T_i * w_i + (1-\alpha) \sum_{i,j} T_i * T_j * v_{ij}$$

subject to

$$\begin{aligned} T_i + T_j \le 1 \quad & \forall (t_i, t_j) \in D \\ T_i - T_j \le 0 \quad & \forall (t_i, t_j) \in H \\ \sum_i T_i \le \delta \end{aligned}$$

where w_i is the weight of the type candidate t_i, α is a hyper parameter, v_{ij} is Pearson correlation coefficient between a type pair (t_i, t_j), D is the set of disjoint type pairs, H is the set of (transitive) hyponym pairs (t_i, t_j) - t_i is the (transitive) hyponym of t_j, and δ is the threshold for the cardinality limit.

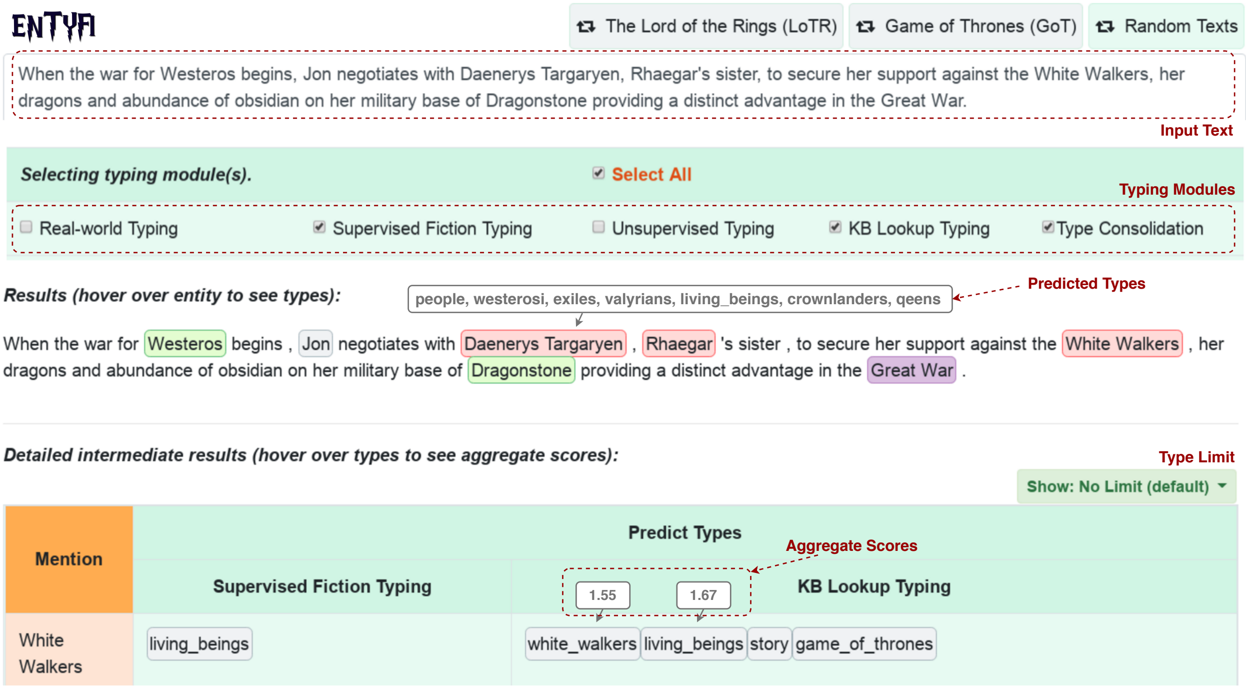

When the war for Westeros begins, Jon negotiates with Daenerys Targaryen, Rhaegar's sister, to secure her support against the White Walkers, her dragons and abundance of obsidian on her military base of Dragonstone providing a distinct advantage in the Great War.

Selecting typing module(s). ☑ Select All

□ Real-world Typing ☑ Supervised Fiction Typing □ Unsupervised Typing ☑ KB Lookup Typing ☑ Type Consolidation

Results (hover over entity to see types): people, westerosi, exiles, valyrians, living_beings, crownlanders, qeens

When the war for Westeros begins , Jon negotiates with Daenerys Targaryen , Rhaegar 's sister , to secure her support against the White Walkers , her dragons and abundance of obsidian on her military base of Dragonstone providing a distinct advantage in the Great War .

Detailed intermediate results (hover over types to see aggregate scores):

Mention	Predict Types	
	Supervised Fiction Typing	KB Lookup Typing
White Walkers	living_beings	white_walkers living_beings story game_of_thrones

Figure 2: ENTYFI Web interface.

3 Web Interface

The ENTYFI system is deployed online at https://d5demos.mpi-inf.mpg.de/entyfi. A screencast video, which demonstrates ENTYFI, is also uploaded at https://youtu.be/g_ESaONagFQ.

Input. The web interface allows users to enter a text as input. To give a better experience, we provide various sample texts from three different sources: Wikia, books and fan fiction[2]. With each source, users can try with either texts from Lord of the Rings and Game of Thrones or random texts, as well as some cross-overs between different universes written by fans.

Output. Given an input text, users can choose different typing modules to run. The output is the input text marked by entity mentions and their predicted types. The system also shows the predicted types with their aggregate scores and the typing modules from which the types are extracted. Figure 2 shows an example input and output of the ENTYFI system.

Typing module selector. ENTYFI includes several typing modules, among which users can choose. If only the real-world typing module is chosen, the system runs typing on the text immediately, using one of the existing typing models which are able to predict up to 112 real-world types (Shimaoka et al., 2017) or 10,331 types (Choi et al., 2018). *Note:* If the later model is selected to run the real-world typing, it requires more time to load the pre-trained embeddings (Pennington et al., 2014).

On the other hand, if supervised fiction typing or KB lookup typing are chosen, the system computes the similarity between the given text and reference universes from the database. With the default option, the type system of the most related universe is being used as targets for typing, while with the alternative case, users can choose different universes and use their type systems as targets. Users are also able to decide whether the consolidation step is executed or not.

Exploration of reference universes. ENTYFI builds on 205 automatically induced high-quality type systems for popular fictional domains. Along with top 5 most relevant universes showing up with similarity scores, users can also choose other universes in the database. For a better overview, with each universe, we provide a short description about the universe and a hyperlink to its Wikia source. Figure 3 show an example of reference universes presented in the demonstration.

Logs. To help users understand how the system works inside, we provide a log box that shows which step is running at the backend, step by step, along with timing information (Figure 4).

[2] https://www.fanfiction.net/

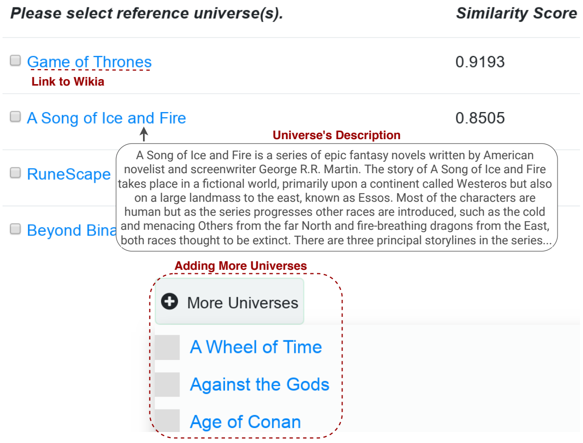

Please select reference universe(s).	Similarity Score
□ Game of Thrones	0.9193
□ A Song of Ice and Fire	0.8505
□ RuneScape	
□ Beyond Bina	

A Song of Ice and Fire is a series of epic fantasy novels written by American novelist and screenwriter George R.R. Martin. The story of A Song of Ice and Fire takes place in a fictional world, primarily upon a continent called Westeros but also on a large landmass to the east, known as Essos. Most of the characters are human but as the series progresses other races are introduced, such as the cold and menacing Others from the far North and fire-breathing dragons from the East, both races thought to be extinct. There are three principal storylines in the series...

Figure 3: ENTYFI Reference Universes.

4 Demonstration Experience

A common use of entity typing is as building block of more comprehensive NLP pipelines that perform tasks such as entity linking, relation extraction or question answering. We envision that ENTYFI could strengthen such pipelines considerably (see also extrinsic evaluation in (Chu et al., 2020)). Yet to illustrate its workings in isolation, in the following, we present a direct expert end-user application of entity typing in fictional texts.

Suppose a literature analyst is doing research on a collection of unfamiliar short stories from fanfiction.net. Their goal is to understand the setting of each story, to answer questions such as what the stories are about (e.g. politics or supernatural), what types of characters the authors create, finding all instances of a type or a combination of types (e.g. female elves) or to do further analysis like if female elves are more frequent than male elves and if there are patterns regarding where female villains appear mostly. Due to time constraints, the analyst cannot read all of stories manually. Instead of that, they can run ENTYFI on each story to extract the entity type system automatically. For instance, to analyze the story *Time Can't Heal Wounds Like These*[3], the analyst would paste the introduction of the story into the web interface of ENTYFI.

"Elladan and Elrohir are captured along with their mother, and in the pits below the unforgiving Redhorn one twin finds his final resting place. In a series of devastating events Imladris loose one of its princes and its lady. But everything is not over yet, and those left behind must lean to cope and fight on."

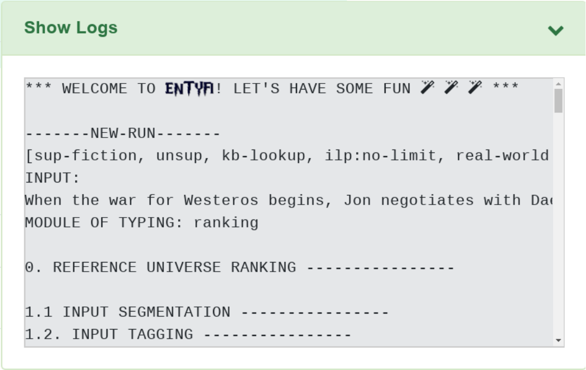

```
*** WELCOME TO ENTYFI! LET'S HAVE SOME FUN ***

-------NEW-RUN-------
[sup-fiction, unsup, kb-lookup, ilp:no-limit, real-world
INPUT:
When the war for Westeros begins, Jon negotiates with Da
MODULE OF TYPING: ranking

0. REFERENCE UNIVERSE RANKING ----------------

1.1 INPUT SEGMENTATION ----------------
1.2. INPUT TAGGING ----------------
```

Figure 4: ENTYFI Logs.

Since they have no prior knowledge on the setting, they could let ENTYFI propose related universes for typing. After computing the similarity between the input and the reference universes from the database, ENTYFI would then propose *The Lord of the Rings, Vampires Diaries, Kid Icarus, Twin Peaks* and *Crossfire* as top 5 reference universes, respectively. The analyst may consider *The Lord of the Rings* and *Vampires Diaries*, top 2 in ranking, of particular interest, and in addition, select the universe *Forgotten Realms*, because that is influential in their literary domain. The analyst would then run ENTYFI with default settings, and get a list of entities with their predicted types as results. They could then see that ELLADAN and ELROHIR are recognized as `living thing, elves, hybrid people` and `characters`, while REDHORN as `living thing, villains, servants of morgoth`, and IMLADRIS as `location, kingdoms, landforms` and `elven cities`.

They could then decide to rerun the analysis with reference universes *The Lord of the Rings* and *Vampires Diaries* but without running type consolidation. By ignoring this module, the number of predicted types for each entity increases. Especially, ELLADAN & EHROHIR now are classified as `living thing, elves, characters`, but also `location` and `organization`. Similarly, REDHORN belongs to both `living thing` and `places`, while IMLADRIS is both a `kingdom` and a `devastating event`. Apparently, these incompatibilities in predictions appear when the system does not run type consolidation.

The analyst may wonder how the system performs when no reference universe is being used. By only selecting the real-world typing module (Choi et al., 2018), the predicted types for EL-

[3] https://www.fanfiction.net/s/13484688/1/Time-Can-t-Heal-Wounds-Like-These

Mention	Settings		
	Default (Ref. universes + all modules)	Default without type consolidation	Only real-world typing
Elladan & Elrohir	men, hybrid peoples, elves of rivendell, real world, elves, characters, living thing, antagonists, supernatural, species, etc.	organization, men, the silmarillion characters, hybrid peoples, elves of rivendell, elves, characters, living thing, location, antagonists, vampire diaries characters, supernatural, etc.	athlete, god, character, body part, arm, person, goddess, companion, brother, child
Redhorn	creatures, villains, servants of morgoth, real world, minions of angmar, servants of sauron, species, living thing, characters, witches, supernatural, one	creatures, villains, evil, death, deaths in battle, servants of morgoth, minions of angmar, servants of sauron, characters, witches, places, arda, races, living thing, organization, etc.	city, god, tribe, county, holiday, body part, society, product, mountain, act
Imladris	kingdoms, location, realms, landforms, places, elven cities, eriador, elven realms, mordor, etc.	kingdoms, location, realms, arda, landforms, places, continents, organization, elven cities, etc.	city, writing, setting, castle, clan, location, character, eleven, etc.

Table 1: Results of ENTYFI on different settings.

LADAN & ELROHIR would change to athlete, god, body part, arm, etc. REDHORN now becomes a city, god, tribe and even an act, while IMLADRIS is a city, writing, setting and castle. The results show not only incompatible predictions, but also that the existing typing model in the real world domain lacks coverage on fictional domains. By using a database of fictional universes as reference, ENTYFI is able to fill these gaps, predict fictional types in a fine-grained level and remove incompatibilities in the final results. From this interaction, the literature analyst could conclude that the story is much related to *The Lord of the Rings*, which might help them to draw parallels and direct further manual investigations. Table 1 shows the result of this demonstration experience in details.

5 Related Work

Earliest approaches for entity typing are based on manually designed patterns (e.g., Hearst patterns (Hearst, 1992)) to extract explicit type candidates in given texts. These pattern-based approaches can achieve good precision, but their recall is low, and they are difficult to scale up.

Traditional named-entity recognition methods used both rule-based and supervised techniques to recognize and assign entity mentions into few coarse classes like person, location and organization (Sang and De Meulder, 2003; Finkel et al., 2005; Collobert et al., 2011; Lample et al., 2016). Recently, fine-grained named-entity recognition and typing are getting more attention (Ling and Weld, 2012; Corro et al., 2015; Shimaoka et al., 2017; Choi et al., 2018). Ling and Weld (2012) use a classic linear classifier to classify the mentions into a set of 112 types. At much larger scale, FINET (Corro et al., 2015) uses 16k types from the WordNet taxonomy as the targets for entity typing. FINET is a combination of pattern-based, mention-based and verb-based extractors to extract both explicit and implicit type candidates for the mentions from the contexts.

With the development of deep learning, many neural methods have been proposed (Dong et al., 2015; Shimaoka et al., 2017; Choi et al., 2018; Xu et al., 2018). Shimaoka et al. (2017) propose a neural network with LSTM and attention mechanisms to encode representations of a mention's contexts. Recently, Choi et al. (2018) use distant supervision to collect a training dataset which includes over 10k types. The model is trained with a multi-task objective function that aims to classify entity mentions into three levels: general (9 types), fine-grained (112 types) and ultra-fine (10201 types).

While most existing methods focus on entity mentions with single contexts (e.g. a sentence), ENTYFI attempts to work on long texts (e.g., a chapter of a book). By proposing a combination of supervised and unsupervised approaches, with a following consolidation step, ENTYFI is able to predict types for entity mentions based on different contexts, without producing incompatibilities in predictions.

Many web demo systems for entity typing have been built, such as Stanford NER[4], displaCy NER[5] and AllenNLP[6]. However, these systems all predict only a few coarse and real world types (4-16 types). ENTYFI is the first attempt to entity typing at a fine-grained level for fictional texts. In a related problem, the richness of Wikia has been utilized for entity linking and question answering (Gao and Cucerzan, 2017; Maqsud et al., 2014).

6 Conclusion

We have presented ENTYFI, an illustrative demonstration system for domain-specific and long-tail typing. We hope ENTYFI will prove useful both to language and cultural research, and to NLP researchers interested in understanding the challenges and opportunities in long-tail typing.

[4] http://nlp.stanford.edu:8080/ner/

[5] https://explosion.ai/demos/displacy-ent

[6] https://demo.allennlp.org/named-entity-recognition

References

Andrew Carlson, Justin Betteridge, Richard C Wang, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Coupled semi-supervised learning for information extraction. In *WSDM*.

Eunsol Choi, Omer Levy, Yejin Choi, and Luke Zettlemoyer. 2018. Ultra-fine entity typing. In *ACL*.

Cuong Xuan Chu, Simon Razniewski, and Gerhard Weikum. 2019. TiFi: Taxonomy induction for fictional domains. In *The Web Conference*.

Cuong Xuan Chu, Simon Razniewski, and Gerhard Weikum. 2020. ENTYFI: Entity typing in fictional texts. In *WSDM*.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. In *JMLR*.

Luciano del Corro, Abdalghani Abujabal, Rainer Gemulla, and Gerhard Weikum. 2015. Finet: Context-aware fine-grained named entity typing. In *ACL*.

Li Dong, Furu Wei, Hong Sun, Ming Zhou, and Ke Xu. 2015. A hybrid neural model for type classification of entity mentions. In *IJCAI*.

Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *ACL*.

Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *NIPS*.

Ning Gao and Silviu Cucerzan. 2017. Entity linking to one thousand knowledge bases. In *ECIR*.

Alex Graves. 2012. Supervised sequence labelling. In *Supervised sequence labelling with recurrent neural networks*.

Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and what's next. In *ACL*.

Marti A Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *COLING*.

Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *NAACL HLT*.

Changki Lee, Yi-Gyu Hwang, Hyo-Jung Oh, Soojong Lim, Jeong Heo, Chung-Hee Lee, Hyeon-Jin Kim, Ji-Hyun Wang, and Myung-Gil Jang. 2006. Fine-grained named entity recognition using conditional random fields for question answering. In *Asia Information Retrieval Symposium*.

Xiao Ling and Daniel S Weld. 2012. Fine-grained entity recognition. In *AAAI*.

Umar Maqsud, Sebastian Arnold, Michael Hülfenhaus, and Alan Akbik. 2014. Nerdle: Topic-specific question answering using wikia seeds. In *COLING*.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.

Marta Recasens, Marie-Catherine de Marneffe, and Christopher Potts. 2013. The life and death of discourse entities: Identifying singleton mentions. In *NAACL*.

Erik F Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. *arXiv*.

Julian Seitner, Christian Bizer, Kai Eckert, Stefano Faralli, Robert Meusel, Heiko Paulheim, and Simone Paolo Ponzetto. 2016. A large database of hypernymy relations extracted from the web. In *LREC*.

Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui, and Sebastian Riedel. 2017. Neural architectures for fine-grained entity type classification. In *EACL*.

Bo Xu, Zheng Luo, Luyang Huang, Bin Liang, Yanghua Xiao, Deqing Yang, and Wei Wang. 2018. Metic: Multi-instance entity typing from corpus. In *CIKM*.

Yu Zhang, Guoguo Chen, Dong Yu, Kaisheng Yaco, Sanjeev Khudanpur, and James Glass. 2016. Highway long short-term memory rnns for distant speech recognition. In *ICASSP*.