



HAL
open science

NanoNER: Named Entity Recognition for nanobiology using experts' knowledge and distant supervision

Ran Cheng, Martin Lentschat, Cyril Labbé

► To cite this version:

Ran Cheng, Martin Lentschat, Cyril Labbé. NanoNER: Named Entity Recognition for nanobiology using experts' knowledge and distant supervision. Proceedings of the 2nd Workshop on Information Extraction from Scientific Publication, 2023, Bali, Indonesia. hal-04423839

HAL Id: hal-04423839

<https://inria.hal.science/hal-04423839v1>

Submitted on 29 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

NanoNER: Named Entity Recognition for nanobiology using experts' knowledge and distant supervision

Ran Cheng, Martin Lentschat, Cyril Labbé

Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, 38000 Grenoble, France
{martin.lentschat, cyril.labbe}@univ-grenoble-alpes.fr

Abstract

Here we present the training and evaluation of NanoNER, a Named Entity Recognition (NER) model for Nanobiology. NER consists in the identification of specific entities in spans of unstructured texts and is often a primary task in Natural Language Processing (NLP) and Information Extraction. The aim of our model is to recognise entities previously identified by domain experts as constituting the essential knowledge of the domain. Relying on ontologies, which provide us with a domain vocabulary and taxonomy, we implemented an iterative process enabling experts to determine the entities relevant to the domain at hand. We then delve into the potential of distant supervision learning in NER, supporting how this method can increase the quantity of annotated data with minimal additional manpower. On our full corpus of 728 full-text nanobiology articles, containing more than 120k entity occurrences, NanoNER obtained a F1-score of 0.98 on the recognition of previously known entities. Our model also demonstrated its ability to discover new entities in the text, with precision scores ranging from 0.77 to 0.81. Ablation experiments further confirmed this and allowed us to assess the dependency of our approach on the external resources. It highlighted the dependency of the approach to the resource, while also confirming its ability to rediscover up to 30% of the ablated terms. This paper details the methodology employed, experimental design, and key findings, providing valuable insights and directions for future related researches on NER in specialized domain. Furthermore, since our approach require minimal man-power, we believe that it can be generalized to other specialized fields.

1 Introduction

As the volume of the scientific literature increases, the demand for NLP models able to deal with domain vocabulary and specific knowledge is becoming

increasingly apparent. The NanoBubbles¹ project, from which the work presented here originates, aims at studying *how, when and why science fails to correct itself*. It focuses on the nanobiology domain and combines approaches from the natural sciences, natural language processing and social sciences. The field of nanobiology being characterized by both its multidisciplinary and its high degree of specialization is a perfect example of the need for specialized tools. Thus, we must leverage methods from Natural Language Processing (NLP) to assist in the extraction of important information from a large number of articles. The main task of this paper is to train a Named Entity Recognition (NER) model in the field of nanobiology.

The primary task of Named Entity Recognition is to identify and classify specific entities (i.e. named entities) in a text. Compared to other fields, Biomedical NER (BMNER) is a particularly challenging problem, mainly due to the high cost of obtaining quality annotated data and the complexity of domain terminology. A famous example of a model able to perform BMNER is bioBERT (Lee et al., 2020), which is pre-trained on a large-scale corpus of biomedical text. It performs well on a standard set of biomedical benchmarks in several downstream tasks (e.g., NER, Relations Extraction, Q&A). To our knowledge, NER in the nanobiology domain remains an uncharted territory, as existing BMNER models are not trained to recognize entities of interest in this specific field.

Training an efficient NER model requires a large amount of annotated data, which is not easy to come by in specialized domains as the manual work it requires need to be carried out by fields experts. In our work for NER in the nanobiology field, we use distant supervision to alleviate for the lack of annotated data and thus allow the creation of a corpus of articles from a specialized domain large enough to train a NER model. Using BioBERT

¹<https://nanobubbles.hypotheses.org/>

(Lee et al., 2020) as base model, this approach requires minimal human work. We believe that the approach we implemented, and describe here, is adapted to other scientific domain.

First, we harnessed existing nanobiology ontologies (i.e., the Nanoparticle ontology (Thomas et al., 2011) and eNanoMapper (Hastings et al., 2015)) for their concept hierarchy and vocabulary. Then, an iterative process took place with a team of domain experts, who determined the essential labels for our NER model and curated the vocabulary. A round of vocabulary extension, with expert curation, took place before the automatic annotation of the corpus. Ablation experiments were also implemented to measure the influence of the vocabulary coverage in our distant supervision setting.

In summary, the main contributions of this paper are as follows:

1. We have implemented a method to create annotated data for NER. It consists in an iterative process, involving ontology and corpus analysis followed by use of expert knowledge and their validation. This lead us to identify five labels, with vocabularies covering 1438 terms, that are highly relevant to nanobiology.
2. We created NanoNER, a NER model for nanobiology using a distant supervision learning approach and trained on automatically annotated entities in a corpus of 728 unlabelled full-text nanobiology articles. Detailed ablation experiments were conducted to evaluate the influence of the vocabulary coverage.
3. Finally, ablation experiments allowed us to estimate the dependency of our model to the annotation resource. We can effectively measure how well NanoNER is capable to generalize, i.e. its ability to (re)find entities not present in the training set, as well as the essential and minimal terms needed to obtain satisfactory results.

2 Related work

Existing BMNER solutions encompass early NER methods, such as dictionary matching or rule-based approaches, as well as supervised machine learning methods such as Markov models (Ponomareva et al., 2007a). Conditional Random Fields (CRFs) were then employed to perform BMNER (Ponomareva et al., 2007b; Friedrich et al., 2006). Unlike

Markov models, CRFs can consider the characteristics of the entire input sequence, not just the current state. And Support Vector Machine (SVM) can be used in binary classification problems for NER tasks, such as determining whether a word is a named entity of a particular type (Ju et al., 2011).

Recently, deep learning approaches using large amounts of labeled data, such as models built on BioBERT (Lee et al., 2020), have achieved state-of-the-art results on BMNER. For instance, on the jnlpba (Huang et al., 2020) dataset, the KeBioLM model (Yuan et al., 2021) obtained a F1 score of 0.82 on recognizing entities relating to proteins, genes and cells. In the bc5cdr (Li et al., 2016) dataset, the BINDER (Zhang et al., 2022) model using a contrastive learning approach, achieved a F1 score of 0.91 on chemical and disease entities. However, BMNER presents specific difficulties. For instance, Dong et al. (2016) conducted an extensive study on electronic medical records and identified that such technical texts often contain a substantial amount of specialized terminology and knowledge, and frequently present issues such as spelling errors, abbreviations, and idiosyncratic terms, all of which add to the difficulty of the NER task. In this difficult setting, they proposed a method based on CNN (Convolutional Neural Networks) and Word2Vec for performing BMNER and managed to achieve a F1 score of 0.73.

To address the scarcity of annotated data in deep learning models, some weak supervision and distant supervision solutions have been proposed. Mintz et al. (2009) were among the pioneers of distant supervision learning, introducing this method in information and relation extraction tasks. Their goal was to extract relations between entities from a large amount of unlabeled text, using existing knowledge bases as distant supervision signals. Distant supervision was initially widely applied to relation extraction tasks and later extensively used in NER tasks. Distant supervision methods for NER have been validated in previous studies. Shang et al. (2018) revised the LSTM-CRF NER model of Lample et al. (2016) and utilized the MeSH database for chemical and disease entity research. Since the automatic annotation of a corpus tend to introduce noise in the training data, some methods have been proposed to reduce this effect (Meng et al., 2021), such as using early stopping or introducing the concept of pseudo-labels (Liang et al., 2020). Early stopping prevent over-

fitting the model on the training data and fosters the learning of important features of the corpus. Pseudo-labeling data expand the training set by generating new labeled data that can then be used alongside existing datasets.

BMNER using ontologies and distant supervision have already been performed in the biomedical domain (Fries et al., 2017; Wang et al., 2021) and this type of approach could be generalized to any domain for which a semantic and lexical resource exists. These works used different techniques to minimize the risk of noise propagation, e.g. filtering candidate annotations through heuristics based on part-of-speech analysis (Fries et al., 2017) or disambiguating ambiguous entities based on other entities present in the same context (Wang et al., 2021). In our work, we rely on domain experts at crucial steps: (1) determining the labels and then (2) filtering and validating the vocabulary of our annotation resource.

To the best of our knowledge, no one has yet proposed a BMNER model that meets the information mining needs of the nanobiology field. We thus aim at training a NER model, using minimal manpower, but which still meets experts requirements regarding the entities of interest of the domain.

3 Data preparation

Here we describe the essential resources for our work, the corpus and ontologies used, the expert work on selecting labels relevant to the domain at hand as well as the vocabulary associated, and the automatic annotation on the scientific articles. All codes necessary to replicate this study are available online².

3.1 Corpus

The corpus used in this study comprises 728 research articles focused on the field of nanobiology. The vast majority of these articles are written in English. In total, the corpus contains 158,283 sentences and 3,762,791 tokens. On average, each paper in the corpus consists of 217 sentences, and each sentence contains approximately 24 tokens. This extensive dataset provides a rich foundation for in-depth analysis and research in the field of nanobiology. The articles were first obtained in PDF format and the abstract and full text of each

²<https://gricad-gitlab.univ-grenoble-alpes.fr/nanobubbles/nano-ner-wiesp-2023.git>

article was extracted using Grobid (Lopez, 2008-2023). Parts of the documents that are not considered as the core of the articles were excluded (e.g. References, Acknowledgment, Appendix).

3.2 Ontology

As resources, we used the NanoParticle Ontology for cancer nanotechnology research (NPO) (Thomas et al., 2011) and eNanoMapper (ENM) (Hastings et al., 2015), which are the two main ontology in the field of nanobiology. As described in the ENM official documentation, ENM is an automatic extension of NPO and reuses several other ontologies including NPO, CHEMINF (Hastings et al., 2011), CHEBI (Degtyarenko et al., 2007) and ENVO (Buttigieg et al., 2013). The NPO possesses 1904 classes and 81 properties, while ENM contains over 25k classes, 697 individuals and 55 properties (August 2023). Since ENM is built automatically we used it as a secondary source to NPO, in order to minimize the risk of noise propagation. The ontologies were used in CSV format, where each concept in the ontology had a unique key, definition, synonyms, and parent key. These resources will be used for their subsumption relations and vocabularies, providing us with a taxonomy and lexical database.

3.3 Labels and vocabulary

To determine the labels our model will be trained to recognize, and their vocabulary, we used an iterative process of reducing the ontologies, expanding the obtained vocabulary and having every steps validated by domain experts. Because of the large number of concepts in the NPO and ENM ontologies, the difficulty of finding a focus to start with and the fact that our aim is to create an automatically labeled corpus, we first retained only the concepts that presented at least one occurrence in our corpus (i.e. $\approx 30\%$ of NPO's and $\approx 10\%$ of ENM's). Concepts that have never appeared in the corpus were discarded, and subsumption relations within the ontology were reconstructed to obtain a reduced ontology.

Using these reduced ontologies, three domain experts (cf. Acknowledgements) examined their structures and the remaining concepts. Together, they identified five labels as being the core concepts of interest to the field of nanobiology, namely Nanoparticle, Property, Material, Event and Technology. Table 1 presents a short description of each label, along

with the core concepts combined under them (the number of their respective sub-concepts before expert selection is indicated between parenthesis) and a vocabulary extract in the last column.

The concepts corresponding to each label are taken as the root concept of ontology sub-trees. We then amalgamated all the terms under the root concepts with all of the terms of all of its respective sub-concepts to build the labels vocabulary. In any conflict between the NPO and ENM structure, NPO was preferred. The labels were subsequently subjected to a first detailed verification by the domain experts, who selected sub-concepts with relevant vocabulary only, which drastically downsized the number (). In addition to verifying each label's vocabulary, they encountered six specific cases of terms under `Material` that they thought should be moved under `Nanoparticle`: *buckyball*, *carbon dot*, *surface group*, *dendrimer*, *liposome* and *fullerene*.

Table 2 presents the characteristics of the labels vocabulary. `Terms` designates the vocabulary size for the label based on the ontologies lexicon. `Vocabulary` indicates the size of the extended vocabulary based on terminological variations retrieval (cf. below), which includes the original terms. `Occurrences` gives the raw frequency of all label's terms in our corpus. Also, since this was obtained by reducing ontologies, the `Depth` and `Width` columns give an insight of the shape of each sub-tree.

After the expert determined the labels and corresponding terms, we recorded the variants of all the terms using FASTR (Jacquemin et al., 1997). Given a list of terms and a corpus of texts, FASTR is able to extract the terminological variations using solely lexical, syntactical and meta-grammatical rules. This tool is also able to account for variations in word order and part-of-speech changes. It can deal with multi-word terms and is able to recognise variations in an expression (e.g. '*molecular function*' → '*functional roles of molecular*'). Although the results of FASTR seemed rather accurate at first, a second round of expert validation of the vocabulary took place. Out of 2,211 unique variations, experts reduced the number to 1,438 terms (i.e. 65%) and thereby preserved the quality of the training data.

3.4 Automatic corpus annotation

We annotated the data for our distant supervision approach using Prodigy (Montani and team, 2023)

under a research licence. The annotation follows the CoNLL2003 (Sang and Meulder, 2003) standard, which uses the BIO annotation format. The Occurrences column in Table 2 displays the number of annotation under each label in our corpus.

4 Experimental methods

The primary objective of our experiments is to test whether the model possesses good generalization capabilities, precision, and stability. Therefore, we designed three distinct ablation studies to evaluate how dependent our approach is to the labels vocabulary.

4.1 Exploring Existing Models

To identify every entity in the articles, we first examined the results of the SciBERT model (Beltagy et al., 2019). SciBERT is a widely pre-trained model for scientific articles, aiming at improving the expressivity of the model and save training time for downstream tasks. We manually annotated 646 "naive" entities (i.e. "naive" meaning only distinguishing whether a span is an entity or not, not knowing which label the entity belongs to) related to the field of nanobiology in one article (Ma et al., 2016), and then tried to use SciBERT for "naive" entity recognition on plain text.

The result is that SciBERT can identify almost all entities in the article. Out of 646 entities related to the nanobiology field it can identify 638, which suggest a high recall capability (i.e. ≈ 0.99 on the article tested). However, SciBERT identifies a large number of entities that would be false positives in the field of nanobiology. SciBERT identified a total of 2,976 entities, which gives 2,322 false positives that need to be filtered out suggesting a low precision value (i.e. ≈ 0.21). Examples of these false positives are : *nanoscience*, *construction*, *convergence*, *reduce*, *Hayakawa* A way to eliminated these false positives would be to match them with an ontology. But this approach would lack several essential aspects: classification of entities into labels, possible confusion between concepts when trying to do so (e.g. in the ontology, *dendrimer* is originally present under the concepts `Material` and `Nanoparticle`), coverage of the ontology vocabulary and so on. Then, it does not eliminate the need for ontology reduction and expert involvement.

We also experimented with some existing models for BMNER in the Scispacy and Stanza li-

| Label | Description | Core Concepts (#sub-concepts) | Vocabulary extract |
|--------------|---|--|--|
| Nanoparticle | are physical structures, usually between 1 and 100nm in two or three dimensions, that present size related properties | Nanoparticle (68), Fiat Material (77) | <i>nanocapsule, fluorescent carbon nanoparticles, carbon dot, nanowire</i> |
| Property | are physical and chemical functions that can be described using measurement | Realizable Entity (89), Application (102) | <i>amphiphilic, hydrophilic, antioxidant, fluorescent</i> |
| Material | are the atoms and chemical compounds constituting nanoparticles and other studied objects | Chemical Entity (663), Material Entity (320) | <i>thiol, gold, primary amine, carbohydrate</i> |
| Event | describes what is happening at a cellular level | Process (231) | <i>mitosis, transcription, cell death, DNA modification</i> |
| Technique | for preparing nanoparticles, measuring their characteristics and using them | Technique (63), Assay (171), Bioassay (37), Instrument (31), Application (102) | <i>fluorescence spectroscopy, atomic force microscopy, gel electrophoresis</i> |

Table 1: Description and examples of the chosen labels

| Label | Terms | Vocabulary | Occurrences | Depth | Width |
|--------------|-------|------------|-------------|-------|-------|
| Nanoparticle | 71 | 196 | 16,341 | 4 | 19 |
| Property | 105 | 345 | 19,849 | 7 | 24 |
| Material | 241 | 515 | 74,688 | 11 | 36 |
| Event | 56 | 210 | 3,219 | 5 | 9 |
| Technique | 65 | 172 | 7,104 | 7 | 19 |
| Total | 538 | 1,438 | 121,201 | | |

Table 2: Labels vocabulary sizes

baries, but most of these are trained on specific corpora and entities, and perform poorly on NER tasks in the nanobiology field.

4.2 Ablation experiment design

In order to assess the dependency of the approach to the resource, as well as model generalization capabilities, we designed a set of ablation experiments. As detailed in Table 2, our five labels cover a list of 538 terms. Each term has varying numbers of variants, ranging from 1 to over 10, resulting in a total of 1438 different terms. In our ablation experiments, the terms were first randomly shuffled in each labels to minimize the risk of latent factors from affecting the experimental results, such as the terms being arranged in a specific pattern. Then,

the label’s vocabularies are divided into five equal parts, noted as folds A, B, C, D and E. Ablation of 33% of the terms were also implemented with folds F, G and H.

To create the training and test set for our ablation experiments, we selected the sentences based on the presence or absence of ablated entities. This was done in order to ensure that the model would not confuse excluded entities for negative examples during the training, and to later test its capabilities to retrieve entities not encountered before. As shown in Table 3, this resulted in training and test sets of various sizes, some being three time larger than others (e.g. test sets in folds D and C). For instance, the training set in Fold A is composed of 126.834 sentences not containing a single ablated

| Fold | Abl.% | Training set | Test set |
|------|-------|--------------|----------|
| A | 20% | 126,834 | 21,427 |
| B | 20% | 130,099 | 18,162 |
| C | 20% | 136,708 | 11,553 |
| D | 20% | 113,875 | 34,386 |
| E | 20% | 129,580 | 18,681 |
| F | 33% | 109,259 | 39,002 |
| G | 33% | 117,392 | 30,869 |
| H | 33% | 120,103 | 28,158 |

Table 3: Number of sentences in each Fold

term, the test set is then made of the remaining sentences. Thus, ablation experiments ensure that part of the test set consist in entities on which the model has not been trained. The approximate sentence ratio is 6:1 for 20% ablation and 4:1 for 33% ablation.

5 Results and Analysis

In this section, we first present the results of training the model on the full dataset, which performances aligned with our expectations. Then, we detail the two random ablation experiments we designed, reducing the data by 20% and 33% respectively. We noticed a significant fluctuation in the results of these experiments. Therefore, we specifically analyzed the 20% ablation experiment and based on these analyses, we proposed the hypothesis that including or excluding specific terms under our labels might have a significant impact on the precision and recall scores. Following this, we designed a new frequency-based 10% ablation experiment to explore this hypothesis. The results from this new experiment successfully validated our conjectures.

5.1 Training on the whole dataset

Initially, we trained the model on the complete dataset, carrying out a training with 5 and 20 epochs respectively. Given the consistency of the training and validation datasets, the nature of this experiment is closer to a straightforward word-matching task. Results are displayed in Table 4. We found that the F1 score of the model reached a value 0.985, which is consistent with our expectations.

Subsequently, we performed a deep analysis of the model’s generalization ability. Our assumption was that it would be impossible to achieve 100% coverage of terms in the corpus, so we had the

model re-annotate the corpus. Table 4 thus also presents the number of unique new entities (i.e. ignoring the number of occurrences) identified by NanoNER, the number of correctly labeled entities and the associated precision. Not considering the number of occurrences for these newly retrieved entities allows for a better estimation of the model generalization capabilities. NanoNER achieves an precision value on new entities roughly around 0.8. Additionally, we found that as the training epochs increased, the precision value on the newly found entities improved, but the number of new entities recognized decreased. This indicates that the number of training epochs can be chosen according to the intended use, giving priority to recall or precision values on never-before-encountered entities.

5.2 20% Ablation experiments

The primary objective of our ablation study is to further test the model generalization capabilities and dependency to the resource. We employed early stopping for training, setting the number of epochs to 1 and the batch size to 32. The training results are presented in Table 5. As we can see, the precision fluctuates around 0.79 and the gap between the highest and lowest precision values can be as high as 0.14. The impact of the vocabulary ablation is even more visible on the recall: with an average score of 0.54, it has degraded considerably compared to the initial training. This tend to indicates that the absence or presence of specific terms highly influence the quality of the trained model. To address this, we conducted further exploration in Section 5.4.1.

5.3 33% Ablation experiments

Next, we conducted experiments with a 33% ablation of the terms. The results in Table 6 are as expected: the precision value remained around 0.8, but the recall rate dropped even further. This is due to the higher number of terms excluded in 33% ablation experiments compared to the 20% ones.

5.4 Ablation experiments analyse

Firstly, we conducted a generalization analysis on the ablation experiments from folds A to E. We employed the same method as in the full data analysis: We initially listed all the deleted entities, and then had predictions made by the model on the entire corpus. Subsequently, we matched all the entities predicted by the model with the deleted entities. As depicted in Table 7, the model could stably

| Epochs | Precision | Recall | F1 Score | New Entities | Correct Labeling | Precision |
|--------|-----------|--------|----------|--------------|------------------|-----------|
| 5.0 | 0.981 | 0.989 | 0.985 | 485 | 375 | 0.773 |
| 20.0 | 0.982 | 0.989 | 0.985 | 249 | 202 | 0.811 |

Table 4: Models trained using the entire dataset

| Fold | Precision | Recall | F1 Score |
|---------|-----------|--------|----------|
| A | 0.796 | 0.544 | 0.646 |
| B | 0.826 | 0.575 | 0.678 |
| C | 0.853 | 0.593 | 0.699 |
| D | 0.756 | 0.424 | 0.544 |
| E | 0.711 | 0.549 | 0.620 |
| Average | 0.788 | 0.537 | 0.637 |

Table 5: Models trained on 20% ablation folds

| Fold | Precision | Recall | F1 Score |
|---------|-----------|--------|----------|
| F | 0.752 | 0.332 | 0.461 |
| G | 0.868 | 0.411 | 0.557 |
| H | 0.847 | 0.460 | 0.596 |
| Average | 0.822 | 0.401 | 0.538 |

Table 6: Models trained on 33% ablation folds

rediscover up to 30% of the deleted entities. Considering that these results are not calculated based on occurrence rates, and it’s quite challenging to discover the relationships between many obscure entities through deep learning, we believe these results are within our expectations.

| Fold | Retrieved | Ablated | Recall |
|---------|-----------|---------|--------|
| A | 66 | 232 | 0.28 |
| B | 89 | 353 | 0.25 |
| C | 72 | 303 | 0.24 |
| D | 77 | 249 | 0.31 |
| E | 77 | 272 | 0.28 |
| Average | 76.2 | 281.8 | 0.27 |

Table 7: Refound

Next, we sought to analyze the variability between the folds. We hypothesize that certain labels are excessively difficult, thereby affecting the overall performance of the task, we then evaluated each label separately. Table 8 displays the average recall and precision values of the labels over the different folds (detailed evaluation is available in Appendix A. We found that *Nanoparticle* and *Event* have the most important variations in recall, while the variations in precision mostly

concern *Nanoparticle* and *Technique*. Most of the average recall values over different labels are close to 0.54, but differences in average precision is more important when comparing the different labels. *Material* and *Property* displays scores over 0.90, *Nanoparticle* is around 0.83, but *Event* and *Technique* have significantly lower precision values (0.74 and 0.65 respectively). These variations are analyzed as a result of the specific characteristics of the labels vocabularies. E.g. *Event* and *Technique* contains terms from the scientific language that are not specific to the nanobiology field, and thus carry a higher risk of confusion.

| Label | Recall | | Precision | |
|---------------------|--------|------|-----------|------|
| | avg. | var. | avg. | var. |
| <i>Nanoparticle</i> | 0.57 | 0.23 | 0.83 | 0.13 |
| <i>Material</i> | 0.55 | 0.13 | 0.92 | 0.06 |
| <i>Event</i> | 0.53 | 0.22 | 0.75 | 0.10 |
| <i>Property</i> | 0.48 | 0.10 | 0.91 | 0.09 |
| <i>Technique</i> | 0.56 | 0.16 | 0.65 | 0.15 |

Table 8: Average performances on each label

We also observed specific differences between the different folds for specific labels. E.g. the recall for *Nanoparticle* is very high in fold A (i.e. 0.82), but significantly lower in fold E (i.e. 0.19). This suggests that certain words are highly important for specific labels. In fact, in fold E, the term *nanoparticle* was removed, which is not only a high-frequency term throughout the entire corpus, but also an essential word involved in different terms (e.g. *nanoparticle*, *gold nanoparticle*).

We believe that these high-frequency terms may have a great impact on the training of the model. Therefore, removing these words during training might lead to a significant decline in the performance of the model. To explore this hypothesis, we decided to conduct a third ablation experiment based on terms of frequency.

5.4.1 10% Ablation experiments

We then sorted the terms according to their frequency in the corpus and conducted four more

| Training set | Precision | Recall | F1 Score | Corpus size |
|--------------------------------|-----------|--------|----------|-------------|
| remove top 10% | 0.599 | 0.279 | 0.381 | 21.2% |
| remove top 10% + <i>m.ft</i> * | 0.668 | 0.371 | 0.430 | 51.0% |
| remove middle 10% | 0.802 | 0.683 | 0.738 | 99.0% |
| top 10% only | 0.763 | 0.413 | 0.536 | 78.8% |

**m.ft*: most frequent term in each label

Table 9: Models trained on 10% ablation

rounds of ablation, as shown in Table 9: in each label we tried removing 10% of the most frequent terms (with one experiment reintroducing the first term in each label), removing the 10% in the middle of the terms frequency and finally retaining only the 10% most frequent terms in each label. This approach limits the corpus exploitable as training and test sets, reflecting the distribution of the terms throughout the corpus.

Firstly, it appears that the frequency of the terms is a good indicator of the model dependency towards the annotation resource. Indeed, precision and recall values are impacted proportionally to the rank of the terms removed. Comparing the first and second rows also indicates that certain terms (i.e. the most frequent terms in each label) significantly impact the model’s performance. These terms may play a critical role in the classification task, or they could provide substantial contextual information, helping the model understand other related terms. Regarding the third row’s ablation experiment, although most of the corpus (99%) was kept to train the model, the F1 score is far lower than 0.985 when trained with the full data. This suggests that even terms with lower frequencies still significantly impact the model’s performance. These less frequent terms might carry specific information crucial for the model to understand and classify the text. The fourth row’s ablation experiment result indicates that retaining only the most common terms might lead the model to overly focus on these terms, overlooking other terms that may carry important information. This could be because these common terms contain a lot of generic information but lack some specific, category-targeted information.

6 Error analysis and improvement approaches

6.1 Sentence selection during ablation experiments

In our ablation experiments, we sometimes encounter a scenario where a sentence contains an

ablated term and an other one that is not, and thus we would like to remove only one of them. In our experiments, we chose to exclude such cases to avoid having our model confuse ablated terms for negative examples. But other strategies could be adopted to tackle this, such as the masking or replacement of tokens.

6.2 Imbalanced dataset

As observed during the ablation experiment, our corpus is highly imbalanced. Some terms and labels appear more frequently than others. Detailed evaluations of individual labels in ablation experiments are given in Appendix A. This discrepancy might lead the model to over-learn from these high-frequency terms, thereby overlooking the importance of less frequent terms. During future optimization, this problem could be tackled by using techniques such as oversampling or undersampling to balance the number of samples across different categories. In a study on NER using Wikipedia, (Al-Rfou et al., 2015) adopted an approach that involves constructing a subset of the training corpus. This strategy ensures that the conditional distribution of specific entity classes remains unaltered when they are positive examples, thereby significantly enhancing the model’s performance across multiple languages.

6.3 Vocabulary coverage in distant supervision

Our training and evaluation assume that the ground truth annotations are accurate, which may not be the case in a distant supervision framework. Thus there may be cases where the model is retrieving entities under the correct label, but that are considered false positives in our automatically annotated corpus. We tried to reduce this effect by employing FASTR (Jacquemin et al., 1997) to improve the coverage of our resource, but this required experts to filter out FASTR false positives. Also, their is some known variations that FASTR is not able to

retrieve (e.g. 'iron oxide nanoparticle' and 'silicon dot' are in `Nanoparticle` vocabulary, but their respective variations 'iron nanoparticle' and 'silicon dot' were not recognized).

One solution would be to use a method known as knowledge distillation, which incrementally improves the model's performance through iterations, a training method used in the previously mentioned BOND (Liang et al., 2020) paper. By using a teacher model to generate pseudo-labels for training the student model, the training effectiveness of the model is improved through repeated iterations. Another solution is to manually annotate a sufficient portion of high-quality data and then use it as validation set for the model.

7 Conclusion

In this work, we have introduced NanoNER, a tool for Named Entity Recognition in the field of nanobiology. We designed an iterative process to determine the model labels and vocabulary using ontologies, domain experts and retrieving terminological variations. This resulted in five labels, covering 1,438 terms, that allow for the automatic annotation of our corpus in a distant supervision approach. Experiment analyses have demonstrated that our model can effectively identify entities of interest, both previously seen and new ones, in the field of nanobiology. Given the complexity and abundance of technical terms in the field, our method shows promising applications in nanobiology.

We believe that this approach can be applied as is on other scientific fields, as it requires only an ontology (or taxonomy) resource and minimum man-power. This allows for the efficient training of NER models useful in downstream NLP tasks.

Ablation experiments showed a significant dependence of the model on the vocabulary used. In future work, we could attempt data augmentation on the dataset to reduce its imbalance and enhance the model's training performance. In addition, it is possible to use knowledge distillation for iterative model updates, which can reduce the false positive misjudgment during validation and improve the model's generalization capabilities.

Acknowledgements

We would especially like to thank the researchers who shared their expertise in nanobiology with us: Raphaël Lévy (Professor of Physics at Paris

Nord University), Nathanne Rost (Postdoctoral researcher in Nanobiosciences at Paris Nord University) and Federico Boem (Postdoctoral researcher in Philosophy of Sciences at University of Twente).

The [NanoBubbles](#) project has received Synergy grant funding from the European Research Council (ERC), within the European Union's Horizon 2020 program, grant agreement no. 951393.

References

- Rami Al-Rfou, Vivek Kulkarni, Bryan Perozzi, and Steven Skiena. 2015. Polyglot-ner: Massive multilingual named entity recognition. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 586–594. SIAM.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [Scibert: A pretrained language model for scientific text](#).
- Pier Luigi Buttigieg, Norman Morrison, Barry Smith, Christopher Mungall, and Suzanna Lewis. 2013. [The environment ontology: contextualising biological and biomedical entities](#). *Journal of biomedical semantics*, 4:43.
- Kirill Degtyarenko, Paula de Matos, Marcus Ennis, Janna Hastings, Martin Zbinden, Alan McNaught, Rafael Alcántara, Michael Darsow, Mickaël Guedj, and Michael Ashburner. 2007. [ChEBI: a database and ontology for chemical entities of biological interest](#). *Nucleic Acids Research*, 36(suppl_1):D344–D350.
- Xishuang Dong, Lijun Qian, Yi Guan, Lei Huang, Qubin Yu, and Jinfeng Yang. 2016. [A multiclass classification method based on deep learning for named entity recognition in electronic medical records](#). In *2016 New York Scientific Data Summit (NYSDDS)*, pages 1–10.
- Christoph M Friedrich, Thomas Revillion, Martin Hoffmann, and Juliane Fluck. 2006. Biomedical and chemical named entity recognition with conditional random fields: The advantage of dictionary features. In *SMBM*.
- Jason Fries, Sen Wu, Alex Ratner, and Christopher Ré. 2017. Swellshark: A generative model for biomedical named entity recognition without labeled data. *arXiv preprint arXiv:1704.06360*.
- Janna Hastings, Leonid Chepelev, Egon Willighagen, Nico Adams, Christoph Steinbeck, and Michel Dumontier. 2011. [The chemical information ontology: Provenance and disambiguation for chemical data on the biological semantic web](#). *PLOS ONE*, 6(10):1–13.
- Janna Hastings, Nina Jeliaskova, Gareth Owen, Georgia Tsiliki, Cristian R Munteanu, Christoph Steinbeck, and Egon Willighagen. 2015. [enanomapper: harnessing ontologies to enable data integration for](#)

- nanomaterial risk assessment. *Journal of biomedical semantics*, 6(1):1–15.
- Ming-Siang Huang, Po-Ting Lai, Pei-Yen Lin, Yu-Ting You, Richard Tzong-Han Tsai, and Wen-Lian Hsu. 2020. [Biomedical named entity recognition and linking datasets: survey and our recent development](#). *Briefings in Bioinformatics*, 21(6):2219–2238.
- Christian Jacquemin, Judith L Klavans, and Evelyne Tzoukermann. 1997. Expansion of multi-word terms for indexing and retrieval using morphology and syntax. In *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pages 24–31.
- Zhenfei Ju, Jian Wang, and Fei Zhu. 2011. Named entity recognition from biomedical text using svm. In *2011 5th international conference on bioinformatics and biomedical engineering*, pages 1–4. IEEE.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *arXiv preprint arXiv:1603.01360*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Jiao Li, Yueping Sun, Robin Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn Mattingly, Thomas Wieggers, and Zhiyong lu. 2016. [Biocreative v cdr task corpus: a resource for chemical disease relation extraction](#). *Database*, 2016:baw068.
- Chen Liang, Yue Yu, Haoming Jiang, Siawpeng Er, Ruijia Wang, Tuo Zhao, and Chao Zhang. 2020. [Bond: Bert-assisted open-domain named entity recognition with distant supervision](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM.
- Patrice Lopez. 2008-2023. [Grobid](#). <https://github.com/kermitt2/grobid>.
- Wang Ma, Andrew G. Cheetham, and Honggang Cui. 2016. [Building nanostructures with drugs](#). *Nano Today*, 11(1):13–30.
- Yu Meng, Yunyi Zhang, Jiaxin Huang, Xuan Wang, Yu Zhang, Heng Ji, and Jiawei Han. 2021. Distantly-supervised named entity recognition with noise-robust learning and language model augmented self-training. *arXiv preprint arXiv:2109.05003*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011.
- Ines Montani and Explosion AI team. 2023. Prodigy - an annotation tool for machine learning and data science. Website. <https://prodi.gy/>.
- Natalia Ponomareva, Ferran Pla, Antonio Molina, and Paolo Rosso. 2007a. Biomedical named entity recognition: a poor knowledge hmm-based approach. In *Natural Language Processing and Information Systems: 12th International Conference on Applications of Natural Language to Information Systems, NLDB 2007, Paris, France, June 27-29, 2007. Proceedings 12*, pages 382–387. Springer.
- Natalia Ponomareva, Paolo Rosso, Ferrán Pla, and Antonio Molina. 2007b. Conditional random fields vs. hidden markov models in a biomedical named entity recognition task. In *Proc. of Int. Conf. Recent Advances in Natural Language Processing, RANLP*, pages 479–483.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the conll-2003 shared task: Language-independent named entity recognition](#).
- Jingbo Shang, Liyuan Liu, Xiang Ren, Xiaotao Gu, Teng Ren, and Jiawei Han. 2018. Learning named entity tagger using domain-specific dictionary. *arXiv preprint arXiv:1809.03599*.
- Dennis G. Thomas, Rohit V. Pappu, and Nathan A. Baker. 2011. [Nanoparticle ontology for cancer nanotechnology research](#). *Journal of Biomedical Informatics*, 44(1):59–74. Ontologies for Clinical and Translational Research.
- Xuan Wang, Vivian Hu, Xiangchen Song, Shweta Garg, Jinfeng Xiao, and Jiawei Han. 2021. Chemner: fine-grained chemistry named entity recognition with ontology-guided distant supervision. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Zheng Yuan, Yijia Liu, Chuanqi Tan, Songfang Huang, and Fei Huang. 2021. Improving biomedical pre-trained language models with knowledge. *arXiv preprint arXiv:2104.10344*.
- Sheng Zhang, Hao Cheng, Jianfeng Gao, and Hoifung Poon. 2022. Optimizing bi-encoder for named entity recognition via contrastive learning. *arXiv preprint arXiv:2208.14565*.

A Detailed ablation evaluation

| Fold | Label | Positive | Total | Recall | Positive | Total | Precision |
|-------------|--------------|-----------------|--------------|---------------|-----------------|--------------|------------------|
| A | nanoparticle | 2,557 | 3,108 | 0.822 | 4,645 | 4,821 | 0.963 |
| | material | 8,952 | 16,022 | 0.557 | 15,433 | 15,821 | 0.975 |
| | event | 233 | 342 | 0.68 | 398 | 444 | 0.896 |
| | property | 2,999 | 8,287 | 0.363 | 4,033 | 4,241 | 0.951 |
| | technique | 591 | 1,036 | 0.57 | 1,346 | 2,053 | 0.656 |
| B | nanoparticle | 1,466 | 3,309 | 0.443 | 2,854 | 4,082 | 0.699 |
| | material | 8,959 | 14,566 | 0.615 | 14,961 | 15,377 | 0.973 |
| | event | 203 | 419 | 0.484 | 341 | 428 | 0.797 |
| | property | 2,531 | 4,122 | 0.613 | 3,964 | 4,101 | 0.967 |
| | technique | 685 | 2,033 | 0.336 | 1,129 | 1,665 | 0.678 |
| C | nanoparticle | 1,346 | 1,900 | 0.708 | 2,475 | 2,590 | 0.956 |
| | material | 6,315 | 9,477 | 0.665 | 10,930 | 11,976 | 0.913 |
| | event | 409 | 1,095 | 0.373 | 554 | 843 | 0.657 |
| | property | 1,385 | 3,410 | 0.406 | 2,229 | 2,295 | 0.971 |
| | technique | 297 | 574 | 0.517 | 958 | 1,240 | 0.773 |
| D | nanoparticle | 2,371 | 3,364 | 0.705 | 4,369 | 4,934 | 0.885 |
| | material | 8,876 | 29,268 | 0.301 | 12,920 | 14,016 | 0.922 |
| | event | 871 | 996 | 0.874 | 1,362 | 2,236 | 0.609 |
| | property | 3,329 | 7,345 | 0.452 | 4,855 | 6,563 | 0.740 |
| | technique | 1,156 | 1,361 | 0.849 | 2,784 | 3,547 | 0.785 |
| E | nanoparticle | 1308 | 6,826 | 0.192 | 2,652 | 4,064 | 0.653 |
| | material | 8,237 | 13,339 | 0.617 | 12,670 | 15,670 | 0.809 |
| | event | 249 | 996 | 0.251 | 366 | 477 | 0.767 |
| | property | 1,749 | 3,056 | 0.573 | 2,921 | 3,251 | 0.898 |
| | technique | 492 | 935 | 0.525 | 1,160 | 3,132 | 0.370 |

Table 10: Ablation evaluation on individual labels