

RESEARCH

Open Access



Parallel sequence tagging for concept recognition

Lenz Furrer^{1,4} , Joseph Cornelius^{2,4} and Fabio Rinaldi^{2,3,4,5*} 

From The 5th workshop on BioNLP Open Shared Tasks Hong Kong, China. 4 November 2019

*Correspondence:

fabio.rinaldi@idsia.ch

² Dalle Molle Institute for Artificial Intelligence Research (IDSIA USI/SUPSI), Lugano, Switzerland

Full list of author information is available at the end of the article

Abstract

Background: Named Entity Recognition (NER) and Normalisation (NEN) are core components of any text-mining system for biomedical texts. In a traditional concept-recognition pipeline, these tasks are combined in a serial way, which is inherently prone to error propagation from NER to NEN. We propose a parallel architecture, where both NER and NEN are modeled as a sequence-labeling task, operating directly on the source text. We examine different harmonisation strategies for merging the predictions of the two classifiers into a single output sequence.

Results: We test our approach on the recent Version 4 of the CRAFT corpus. In all 20 annotation sets of the concept-annotation task, our system outperforms the pipeline system reported as a baseline in the CRAFT shared task, a competition of the BioNLP Open Shared Tasks 2019. We further refine the systems from the shared task by optimising the harmonisation strategy separately for each annotation set.

Conclusions: Our analysis shows that the strengths of the two classifiers can be combined in a fruitful way. However, prediction harmonisation requires individual calibration on a development set for each annotation set. This allows achieving a good trade-off between established knowledge (training set) and novel information (unseen concepts).

Keywords: Text mining, Named entity recognition and normalization, Concept recognition, Neural network, Sequence tagging

Background

Concept recognition is a fundamental task in text mining for biomedical texts. Biomedical text mining finds applications in literature analysis, literature-based discovery but also over other types of text, such as clinical records and social media. For most applications, identifying occurrences of biomedical concepts is an essential first step. The task is usually tackled in a two-stage approach: First, named entity recognition (NER), or span detection, is concerned with identifying textual mentions of relevant entities, such as proteins, chemicals, or species. Second, the identified mentions are assigned to a concept entry in a controlled vocabulary, which is referred to as named entity normalisation



(NEN), linking, or grounding. Typically, the two steps are performed in a sequential manner, using a sequence classifier for NER and a ranking- or rule-based module for NEN. While this approach allows focusing on different methods for the individual steps, it suffers from error propagation, an inherent drawback of any pipeline architecture. For example, a certain NEN system might have excellent accuracy when using ground-truth spans as input, but its performance will decrease when operating on the imperfect output of a span tagger. In particular, a normaliser might be inclined or even forced to predict a concept ID for spurious spans, and it cannot recover from cases where a span is missing.

In this work, we investigate an alternative architecture for concept recognition, which alleviates the problem of error propagation: parallel sequence tagging for NER and NEN. In this architecture, NEN is modeled as a sequence-classification problem (like NER) and applied to the input text independently of the span tagger. The predictions of the two taggers are harmonised using different strategies, the choice of which is a hyperparameter of the complete system. We test our approach with a manually annotated dataset for biomedical concepts, the CRAFT corpus, continuing the efforts from our participation in the CRAFT shared task 2019.

Related Work

Concept recognition has often been approached as a pipeline of NER+NEN. For NER, sequence labeling with conditional random fields (CRF) has dominated the field to present, be it pure CRF as in Gimli [1] or DTMiner [2], on top of a recurrent neural network as in HUNER [3], Saber [4], or DTranNER [5], or even as the head of a BERT-based system as in SciBERT [6]. BERN [7] performs NER by fine-tuning BioBERT alone, even though [8] report improved results when stacking CRF atop BioBERT. Different approaches have been taken to NEN, where extracted mentions are mapped to a vocabulary: exact match as in Neji [9], expert-written rules [10], learning-to-rank as in DNORM [11], linking through an ontology using word embeddings and syntactic re-ranking [12], or sequence-to-sequence prediction [13].

Knowledge-based concept-recognition systems like Jensen tagger [14] or NOBLE coder [15] do not allow for a clear separation between NER and NEN, as span detection and linking happens at once, even if machine-learning components are added for improving accuracy, like for OGER++ [16] or RysannMD [17]. Joint approaches like TaggerOne [18], JLink [19], and others [20, 21], however, have separate modules for NER and NEN, which are trained simultaneously. The multi-task sequence labeling architecture for NER and NEN in [21] has been highly inspirational for the present work, although we were unable to reproduce their results, even using the code that the authors made publicly available.

CRAFT corpus and shared task

The Colorado Richly Annotated Full-Text (CRAFT) corpus [22, 23] is a collection of 97 scientific articles from the biomedical domain. It is manually annotated for syntactic structure, coreferences, and bio-concepts (entities), the last of which are used in the present study. In the latest release (Version 4), the concept annotations are divided into 10 sets of different entity types, which are provided in two versions each (proper and

extended¹), for a total of 20 separate annotation sets over the same text collection. The concepts are linked to 8 different ontologies, as shown below (ontology in parentheses):

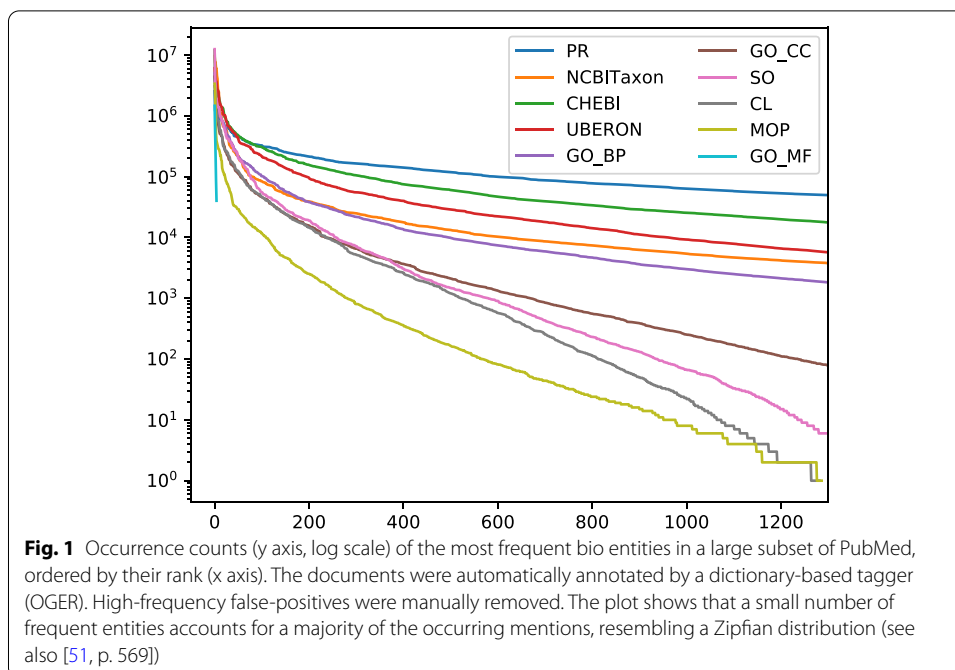
- CHEBI:** chemicals/small molecules (Chemical Entities of Biological Interest [24])
- CL:** cell types (Cell Ontology [25])
- GO_CC:** cellular and extracellular components and regions (Gene Ontology [26])
- GO_BP:** biological processes (Gene Ontology)
- GO_MF:** molecular functionalities possessed by genes (Gene Ontology)
- MOP:** chemical reactions and other molecular processes (Molecular Process Ontology [27])
- NCBITaxon:** biological taxa and organisms (NCBI Taxonomy [28])
- PR:** proteins, genes, and transcripts (Protein Ontology [29])
- SO:** biomacromolecular entities, sequence features (Sequence Ontology [30])
- UBERON:** anatomical entities (UBERON [31])

The extended annotations are referred to by appending EXT to the abbreviations for the proper annotations (CHEBI_EXT, CL_EXT etc.).

The CRAFT corpus has been used in a range of studies. Through repeated improvements and extensions over time, the corpus has become a high-quality resource with rich annotations, but it also led to the situation that most experiments are not directly comparable to each other, as their setup differs in many ways. In the first release of the CRAFT corpus, only 67 articles were available. The remaining 30 documents were not released until the evaluation period of the CRAFT shared task 2019 [32], where they served as a test set. This competition was part of the BioNLP Open Shared Tasks and comprised three core NLP tasks, where participating systems were evaluated against the ground-truth annotations of Version 4 of the CRAFT corpus. However, most prior work on concept recognition was carried out with an older version of CRAFT, i. e. using a different test set, possibly an earlier stage of annotations and a different evaluation method, which means that results are not directly comparable.

While the majority of studies is concerned with concept recognition (i. e. systems that predict IDs), some are restricted to NER, e. g. [4, 33, 34]. Methodologically, the approaches range from pure dictionary-based [15, 35] to entirely example-based systems [36], even though the NER step almost always includes dictionary lookup. Since no official test set was available prior to Version 4, many experiments use an arbitrary train/test split [37] or apply evaluation to the entire corpus [9]. The metrics used are consistently precision, recall and F-score, but differences exist with respect to considering partial matches. Also, many studies do not cover the full set of annotations, but rather focus on a small selection of entity types, such as Gene Ontology [38] or gene mentions [33].

¹ The extended annotations are based on a modified version of the reference ontologies. Through these modifications, the corpus creators aimed at more accurately capturing language use in scientific literature.



Methods

We propose a paradigm for biomedical concept recognition where named entity recognition (NER) and normalisation (NEN) are tackled in parallel. In a traditional NER+NEN pipeline, the NEN module is restricted to predict concept labels (IDs) for the spans identified by the NER tagger. In order to avoid the error propagation inherent to this serial approach, we drop this restriction and provide the full input sequence to the normaliser. As such, we cast the normalisation task as a sequence-tagging problem – very much like an NER tagger, but with a considerably larger tag set, consisting of all concept IDs of the training data.

Design implications

Modeling concept normalisation as sequence tagging has a number of drawbacks. As discussed in the next section, the CoNLL representation of the data enforces exactly one label for each token, which disallows learning and predicting annotations with overlapping and discontinuous spans. This representation also entails that the model has to produce a consistent series of individual predictions in order to correctly label a multi-word expression. This often means that highly ambiguous tokens like prepositions, numbers, or single letters must be interpreted correctly in context (e. g. “of” in “inhibitor of calpain”, “I” in “hexokinase I”). As the most serious limitation, a sequence tagger can only ever predict labels it has seen during training, which restricts the label set of the trained system to a fraction of the target label set (the ontology) in many cases. Since many concepts occur extremely rarely in the biomedical literature (cf. Fig. 1), this limitation might not critically reduce performance measured on a typical evaluation data set. However, it is highly undesirable to have a tagger that is completely incapable of predicting labels beyond the training set.

On the other hand, the ID-tagging architecture is technically an end-to-end concept-recognition system, i. e. it does not depend on any span predictions, which means that the NER step could potentially be skipped entirely. However, due to the small number of tags, span tagging is far more robust with respect to ambiguous tokens and unseen concepts. By adding span predictions, we might thus be able to overcome the limitations of direct ID tagging. Therefore, we chose to combine the strengths of span and ID tagging by applying both in parallel and merging the results in postprocessing.

Data preparation

Our system processes documents in a variant of the CoNLL format, i. e. a verticalised format where each text token is assigned exactly one label. Based on our architecture with two sequence classifiers, we employed two different label sets. For the span tagger, the text is tagged with IOBES labels, i. e. each token is assigned one of the five labels I, O, B, E, or S. Entities spanning only a single token are annotated with S. For multi-word entities, the first and last token are tagged with B and E, respectively, and any intervening tokens with I. The rest of the text (i. e. all tokens outside of an entity) are annotated with O. For the ID tagger, all tokens of an entity are tagged with the respective concept ID. We added a NIL label to mark non-entity tokens, analogously to the O tag of the span tagger.

This representation does not have the same expressiveness as the stand-off format used in CRAFT, which offers great flexibility for anchoring annotations in the text. In particular, the CRAFT corpus contains discontinuous annotations (multiple non-adjacent text spans for the same annotation), overlapping annotations (words shared by multiple annotations) and sub-word spans (annotation refers to part of a word). Since these complex annotations cannot be represented with token-level labels, their structure needs to be simplified.

In order to measure the performance impact of this simplification, we converted the reference annotations of the training set to CoNLL format and back to stand-off using the *standoff2conll* suite [39]. This utility offers two strategies for unifying discontinuous annotations (*full-span* and *last-span*), to which we added a third option (*first-span*) [40]. For unnesting overlapping annotations, two strategies are available as well (*keep-longer* and *keep-shorter*). The effect of unifying and unnesting annotations is illustrated in Fig. 2. Sub-word annotations are extended to span entire tokens.

After this round-trip conversion, the annotations are run through the official evaluation suite provided by CRAFT [41]. Table 1 shows the results for different combinations of unification and unnesting strategies on the non-extended annotation sets. These numbers mark the upper limit for a system trained on input data in CoNLL format. For all annotation sets, using the *first-span* and *keep-longer* strategies achieved the highest F-score.

Architecture

The sequence taggers used in our experiments are built atop a pretrained language-representation model, BioBERT [42], which in turn extends BERT [43]. BERT is an attention-based multi-layer neural network which learns context-dependent word-vector representations. It creates bidirectional contextual representations of a token from

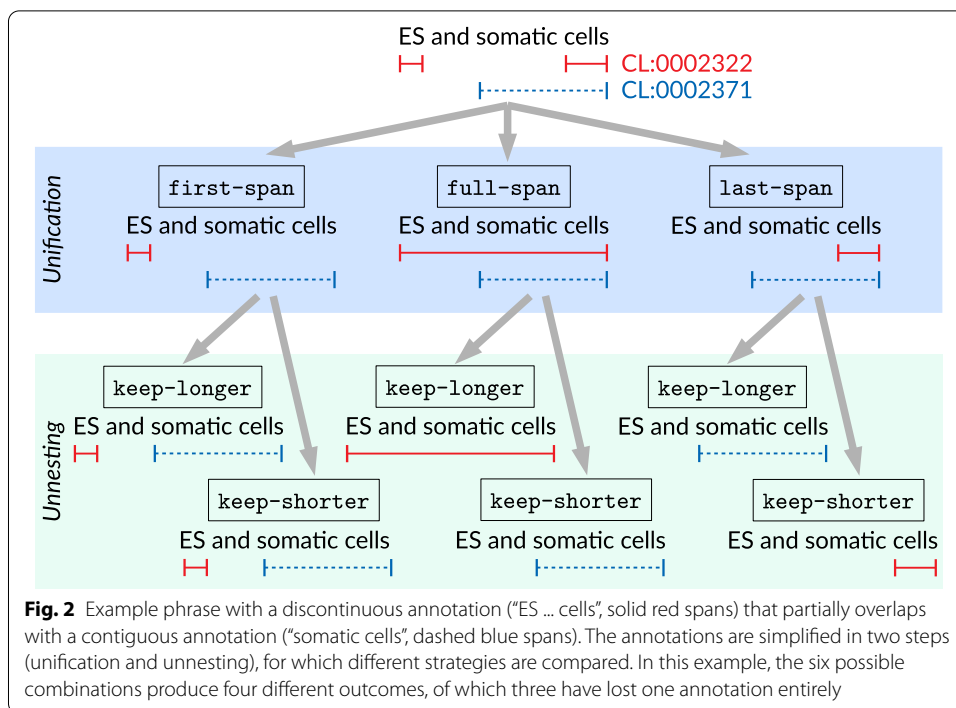


Table 1 Upper bound of annotation performance (F-score) when using the CoNLL format, comparing different simplification strategies

Unification	First-span		Full-span	Last-span
Unnesting	Keep-shorter	Keep-longer	Keep-longer	
CHEBI	0.9979	0.9980	0.9972	0.9974
CL	0.9706	0.9720	0.9574	0.9692
GO_BP	0.9607	0.9626	0.9587	0.9570
GO_CC	0.9811	0.9813	0.9801	0.9785
GO_MF	0.9974	0.9974	0.9974	0.9974
MOP	0.9967	0.9967	0.9967	0.9967
NCBITaxon	0.9996	0.9996	0.9995	0.9996
PR	0.9624	0.9627	0.9619	0.9618
SO	0.9829	0.9831	0.9816	0.9825
UBERON	0.9792	0.9798	0.9776	0.9780

For each annotation type (row), the best value is given in bold

For reasons of clarity, the combinations *full-span/keep-shorter* and *last-span/keep-shorter* are omitted; their results are in most cases inferior to those presented in this table

unlabeled text conditioned on the left and the right context. BERT is trained to solve two tasks: first, to predict whether two sentences follow each other, and second, to predict a randomly masked token from its context. After a slight modification to its architecture, training of BERT can be continued on a different task like NER; this process is referred to as fine-tuning with a task-specific head.

For our experiments, we downloaded BioBERT v1.1, which includes code, configuration and pretrained parameters. BioBERT is based on BERT_{BASE}, which was pretrained for 1M steps by Devlin et al. [43] on a 3.3B-word corpus from the general domain

(English Wikipedia, BooksCorpus). Lee et al. [42] continued training for another 1M steps on a 4.5B-word biomedical corpus (PubMed abstracts). Finally, we fine-tuned BioBERT for sequence-tagging on the CRAFT corpus for 55 epochs (approximately 53k steps).

To perform NER and NEN in parallel, we used two different tag sets for fine-tuning, as described in the previous section: IOBES labels for the span tagger and the set of all concept IDs for the ID tagger. In addition to that, both taggers used a small set of tags inherited from the original BERT implementation, which flag tokens with a special function, such as padding, sub-word unit and sentence boundary. We trained a pair of span and ID tagger for each annotation set, which resulted in a total of 40 individual models.

The predictions of the span tagger are always aligned with the IDs produced by a dictionary-based concept-recognition system, OGER [16, 44]. OGER detects mentions of ontology terms in running text through efficient fuzzy-matching. We manually optimised OGER's configuration on the CRAFT training set. We used no additional terminology resources besides the ontologies provided with the corpus. However, we manually added a handful of synonyms for GO_MF. This combined system resembles a classical NER+NEN pipeline, where the high-recall output of the dictionary-based system is combined with the context-aware span detection using an example based classification model.

Hyperparameter tuning

In order to determine the best hyperparameters for each annotation set, we performed extensive grid search in cross-validation over the training set. In particular, we investigated the following configurations:

- ontology pretraining:** enable/disable
- abbreviation expansion:** enable/disable
- prediction harmonisation:** 6 strategies

If ontology pretraining is enabled, the ID classifier is trained on synonym-ID pairs from the terminology for 20 epochs before switching to the actual training corpus. For abbreviation expansion, we first used Ab3P [45] to detect abbreviation definitions, then replaced occurrences of short forms with the corresponding long form. For harmonising the predictions of the two classifiers, we compared six different strategies; these are described in the next section.

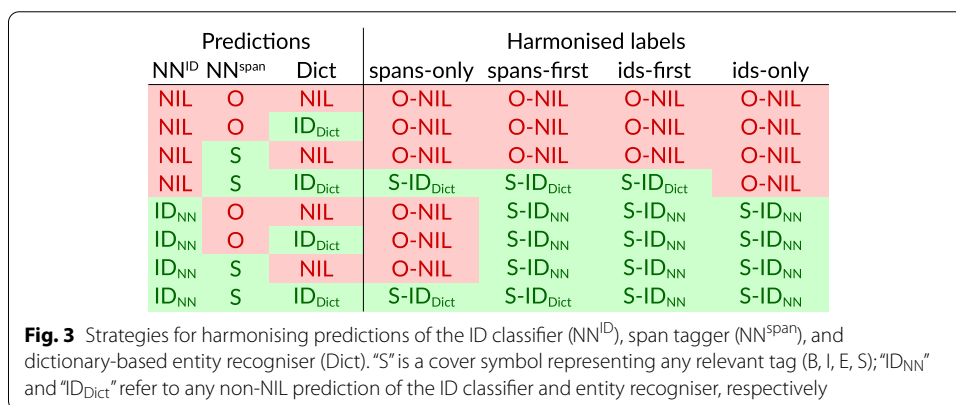
From previous experiments [46], we knew that ontology pretraining has a positive effect for some, but a negative effect for other annotation sets. We therefore concluded that hyperparameters had to be tuned individually for each of the 20 annotation sets. In order to obtain reliable figures, we performed 6-fold cross-validation with up to 3 runs for each combination.

As we expected, ontology pretraining yielded a mixed picture. In many cases, a clear decision was not possible, as repeated runs gave contradictory results. Unexpectedly, abbreviation expansion showed a clear improvement only for CL and a slight improvement for GO_MF; in all other cases (including CL_EXT and GO_MF_EXT)

Table 2 Best-performing harmonisation strategy by annotation set, based on 6-fold cross-validation over the training set

Annotation set	Proper	Extended
CHEBI	spans-first	spans-first
CL	spans-first	ids-only
GO_BP	spans-first	ids-only
GO_CC	spans-first	spans-first
GO_MF	spans-/ids-first	spans-first
MOP	spans-/ids-first	ids-first
NCBITaxon	ids-first	ids-first
PR	spans-only	spans-only
SO	ids-only	ids-only
UBERON	spans-first	spans-first

For GO_MF and MOP, results for *spans-first* and *ids-first* are identical



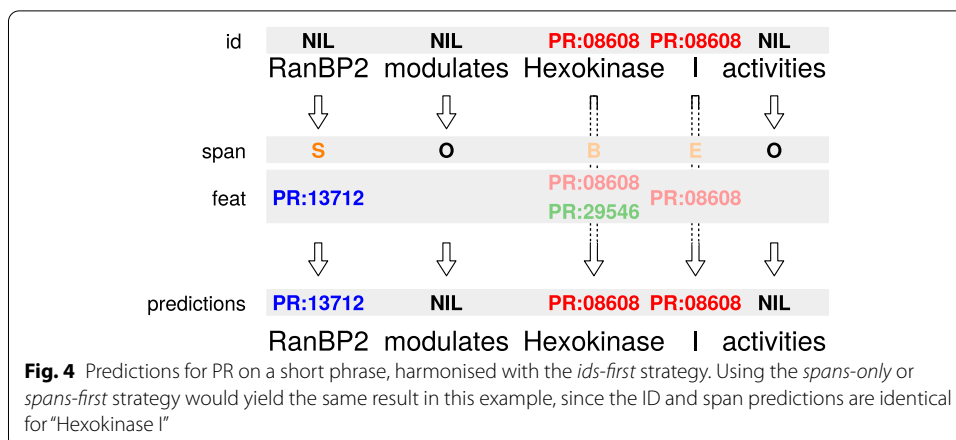
the results decreased. We decided to disable both ontology pretraining and abbreviation expansion, as the isolated merits do not justify the added complexity.

For prediction harmonisation, the best strategy for each annotation set is given in Table 2 and discussed in the following section. The full results for the whole tuning phase are included in Additional file 1.

Harmonising predictions

The predictions of the span and ID classifier are not guaranteed to agree, even if trained jointly. Disagreement occurs if the span classifier predicts a relevant tag (B, I, E, S) for a particular token while the ID classifier predicts NIL, or, conversely, if the ID classifier predicts a specific concept for a token tagged as irrelevant (O) by the span classifier. In addition, the dictionary feature of the knowledge-based entity recogniser might or might not agree with the neural predictions. This results in $2 \times 2 \times 2 = 8$ prediction patterns concerning the relevance of a given token.

We considered four different strategies for harmonising conflicting predictions: *spans-only*, *ids-only*, *spans-first*, and *ids-first* (cf. Fig. 3). These strategies are heuristics with a predetermined bias towards one of the two classifiers. Two additional strategies (*mutual* and *override*), which use the confidence scores for balancing the classifiers, consistently



produced worse results compared to the simpler bias strategies. The score-based strategies are thus not discussed here; however, we used and described the *mutual* strategy when participating in the CRAFT shared task [46, p. 188]. The systematic application of different harmonisation strategies is one of the major differences of this work compared to the work presented at the shared-task workshop.

With the *spans-only* strategy, the ID predictions are completely ignored. In order to provide a concept label, the span predictions are combined with the dictionary feature provided by OGER; in case of multiple features, an arbitrary decision is taken (lexically lowest ID). Since a concept label is always required, span predictions without a supporting feature have to be dropped.

With the *ids-only* strategy, the predictions are based primarily on the ID predictions, whereas the span predictions are overridden (e. g. the span tag cannot be O when the ID classifier predicts a non-NIL concept). The dictionary feature is ignored in the decision.

The *spans-first* and *ids-first* strategies are combinations of the previous two. With the former, the *spans-only* strategy is applied first, backing off to the *ids-only* strategy if the outcome is O-NIL. Analogously, the *ids-first* strategy gives preference to *ids-only*. An example with partially disagreeing predictions is given in Figure 4.

We compared the effect of the different strategies in a 6-fold cross-validation over the training set. For each annotation set, we determined the best harmonisation strategy based on F-score according to the official evaluation suite. As shown in Table 2, using both span and ID predictions was beneficiary most of the time. In many cases, the same strategy worked best for the proper and extended classes. Intuitively, the choice of *spans-only* for proteins makes sense, as PR[_EXT] shows an exceedingly high number of different concepts with a small overlap between training and test data, which is a tough scenario for the ID tagger. Conversely, entity types with a limited number of distinct concepts in the corpus like sequences and organisms rely more heavily on the ID tagger. The choice of harmonisation strategy was fixed as a hyperparameter for the test-set predictions.

Results and discussion

We evaluated our concept-recognition system using the official evaluation suite [41]. Performance is measured in terms of F-score, i. e. the harmonic mean of precision and recall, and slot error rate (SER) [47]. Both metrics are based on the counts of matches (true positives), substitutions (partial errors), insertions (false positives), and deletions (false negatives). Partially correct predictions are assigned a similarity score m in the range $[0, 1]$, which measures the accurateness of the predicted spans and concept labels [48]. The similarity score incorporates a notion of textual overlap (Jaccard index at the character level) and a weighted measure of shared ancestors in the ontology hierarchy, as introduced in [49]. The fractional value m is added to the match count, whereas the remainder $1 - m$ is counted as a substitution. While precision, recall, and F-score are figures of merit ranging from 0 (worst) to 1 (best), SER is a measure of error that assigns 0 to a perfect system and higher values to lower performance. Even though the values for SER and F-score often correlate, they are not guaranteed to produce identical rankings. In particular, SER is more sensitive to false-positive errors than F-score, and low precision has a stronger impact on SER than low recall. Please note that perfect scores cannot be reached by our systems due to limitations in the input representation, as explained in the *Data preparation* section.

The results for our parallel NER+NEN system are given in Table 3. The scores are compared to our systems developed for the shared task [46] and to the official baseline published in the workshop overview [32]. Our system consistently achieves better scores than the baseline, which is a pipeline with a CRF-based span tagger and a BiLSTM-based concept classifier that were also trained on the CRAFT corpus alone. For most annotation sets, our current system performed better than the best system presented in the shared-task paper, with the exception of GO_MF_EXT and PR_EXT. For NCBITaxon_EXT and PR, the comparison is inconclusive, as SER and F-score give contradictory rankings.

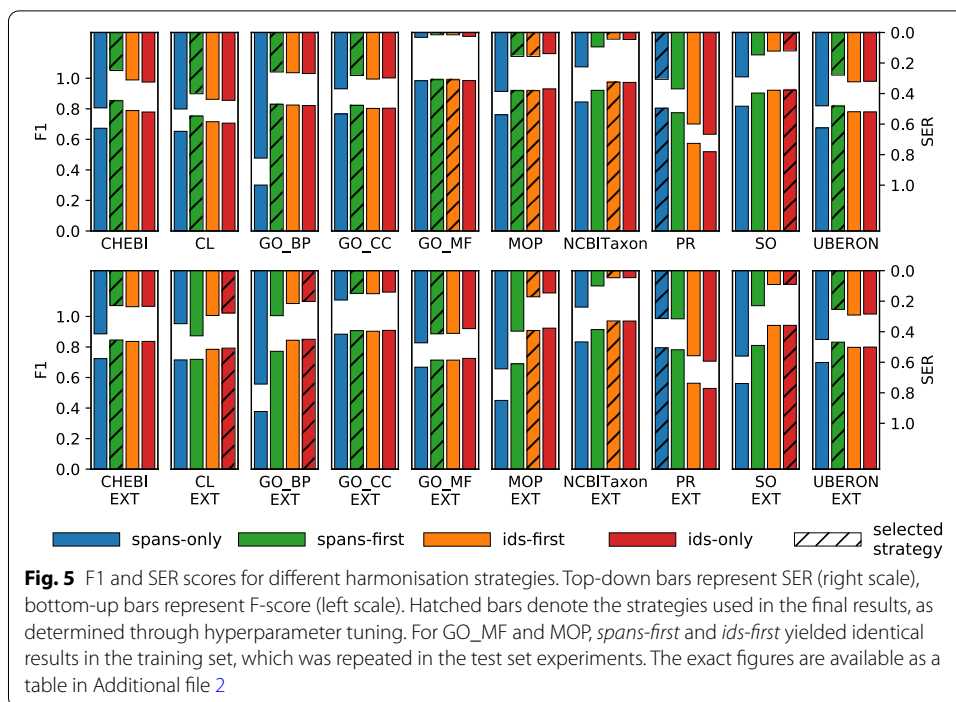
Unfortunately comparison with other systems is difficult due to the fact that the complete CRAFT corpus was not available before the shared task. Previous published results on the CRAFT corpus (such as [50]) are based on a different (and smaller) version of the corpus.

Effect of harmonisation

In order to measure the effect of the different harmonisation strategies, we evaluated all four strategies on the test set, as shown in Fig. 5. This study also serves as a validation for our hyperparameter-tuning approach, i. e. whether cross-validation on the training set can be used for reliably picking the best-suited harmonisation strategy. For the majority of the annotation sets, the picked strategy also worked best for the test set. Where the picked strategy was not the best (GO_MF_EXT, MOP[_EXT]), the difference to the top-performing strategy was comparatively small.

Unseen concepts

As stated above, a major limitation of trained sequence labeling for IDs is the inability to predict concepts not seen among the training examples. An important goal of combining



the ID tagger with a span tagger and dictionary-based predictions is to overcome this limitation. To study the effect of the different harmonisation strategies on unseen concepts, we performed another evaluation on a subset of the annotations. To this end, we filtered both ground truth and predictions of the test set to contain only annotations with concept labels that are not used in the training set.

Table 4 shows precision and recall scores as well as annotation counts for the subset of unseen concepts. The *ids-only* strategy is omitted in the table, as this configuration can never predict unseen concepts. The *spans-only* and *spans-first* strategies systematically yield identical results, as they only differ in cases where the latter backs off to ID predictions, which have been filtered out in this evaluation. With the *ids-first* strategy, many span predictions for unseen concepts are shadowed by an ID prediction for a concept known from the training set (which is then ignored in this specific evaluation). For some annotation types (e. g. CHEBI[_EXT], GO_BP[_EXT], SO[_EXT]), the removal of known concepts improves precision, i. e. more false positives than true positives were removed. In other cases, precision suffers from the removal. Recall decreases in all cases, as is to be expected for an evaluation that focuses on more difficult examples.

Interpretation

Tackling concept recognition for multiple entity types with a single architecture is very challenging, even if a separate model is trained for every annotation set. The comparative results for the different harmonisation strategies (Figure 5) illustrate well how some annotation sets profit more from the span tagger (blue, left-most bars), others more from the ID tagger (red, right-most bars). In many cases, merging predictions from the two taggers (middle bars) yields better results than relying on a single tagger (outer

Table 3 Results for our current BioBERT system, best system reported in the shared-task paper [46], and the official baseline

Annot. set	System	Proper		Extended	
		SER	F-score	SER	F-score
CHEBI	Baseline	0.44	0.72	0.29	0.80
	Shared-task	0.3388	0.7700	0.2571	0.8209
	Current	0.2492	0.8528	0.2289	0.8459
CL	Baseline	0.53	0.61	0.33	0.73
	Shared-task	0.4862	0.6657	0.3361	0.7484
	Current	0.4013	0.7526	0.2777	0.7926
GO_BP	Baseline	0.39	0.72	0.29	0.79
	Shared-task	0.3047	0.8037	0.2786	0.8138
	Current	0.2587	0.8297	0.2015	0.8506
GO_CC	Baseline	0.44	0.71	0.20	0.88
	Shared-task	0.3788	0.7645	0.1678	0.8936
	Current	0.2817	0.8219	0.1486	0.9073
GO_MF	Baseline	0.07	0.95	0.45	0.66
	Shared-task	0.0319	0.9838	0.3881	0.7438
	Current	0.0149	0.9904	0.4135	0.7139
MOP	Baseline	0.43	0.75	0.36	0.79
	Shared-task	0.2684	0.8705	0.3080	0.8437
	Current	0.1567	0.9188	0.1713	0.9082
NCBITaxon	Baseline	0.07	0.96	0.07	0.96
	Shared-task	0.0537	0.9694	0.0466	0.9722
	Current	0.0436	0.9744	0.0460	0.9704
PR	Baseline	0.69	0.48	0.62	0.52
	Shared-task	0.3052	0.8026	0.3030	0.8011
	Current	0.3068	0.8041	0.3130	0.7951
SO	Baseline	0.21	0.86	0.18	0.89
	Shared-task	0.1593	0.9027	0.1230	0.9187
	Current	0.1206	0.9223	0.0899	0.9419
UBERON	Baseline	0.41	0.70	0.36	0.75
	Shared-task	0.3752	0.7488	0.3371	0.7714
	Current	0.2790	0.8177	0.2537	0.8315

In case of the shared-task systems, the results were selected independently for SER and F-score, i. e. the two scores for a given annotation set do not necessarily come from the same system. For the baseline and the current BioBERT system, however, only one system was evaluated for each annotation set

bars). This preference does not directly correlate with ontology size: the two annotation sets with the largest ontologies (NCBITaxon and PR) show quite distinct result patterns. However, it is possible to empirically determine how well each harmonisation strategy suits the characteristics of a given annotation set. Using cross-validation over the training set resulted in robust estimations for ranking the harmonisation strategies.

The diversity of the individual annotation sets shows even more clearly when it comes to predicting unseen concepts. In general, the level of precision and recall for unseen concepts varies greatly across annotation sets, as does the number of unseen concepts in the reference (cf. Table 4). There is a loose negative correlation to the performance on the entire test set: annotation sets like NCBITaxon[_EXT] and SO[_EXT] show high overall scores and low scores for unseen concepts, whereas more

Table 4 Precision and recall for unseen concepts in the test set

Annotation set	Harmonisation	Proper				Extended			
		Ref. count		Prec.	Recall	Ref. count		Prec.	Recall
		Unique	Occ.			Unique	Occ.		
CHEBI	Spans-only/-first	110	447	0.7747	0.5199	134	538	0.6265	0.5462
	lds-first			0.8805	0.0867			0.7131	0.0530
CL	Spans-only/-first	52	484	0.8817	0.2222	52	484	0.6900	0.2338
	lds-first			0.8750	0.0723			0.4326	0.0152
GO_BP	Spans-only/-first	120	484	0.6170	0.1402	126	508	0.2699	0.1621
	lds-first			0.7466	0.0524			0.4243	0.0175
GO_CC	Spans-only/-first	32	184	0.6343	0.1965	36	231	0.5158	0.1853
	lds-first			0.4956	0.0458			0.3375	0.0058
GO_MF	Spans-only/-first	1	1	–	–	73	416	0.5366	0.1393
	lds-first			–	–			0.5329	0.0090
MOP	Spans-only/-first	2	2	–	–	2	2	–	–
	lds-first			–	–			–	–
NCBITaxon	Spans-only/-first	40	87	0.3805	0.2974	44	95	0.4070	0.3342
	lds-first			0.7346	0.1689			0.7363	0.1163
PR	Spans-only/-first	278	4782	0.8170	0.7350	309	5156	0.8286	0.7230
	lds-first			0.7402	0.1200			0.6852	0.0909
SO	Spans-only/-first	16	101	0.0962	0.0571	25	123	0.1586	0.2579
	lds-first			1.0000	0.0198			0.9345	0.1823
UBERON	Spans-only/-first	203	1297	0.7246	0.2447	207	1308	0.7142	0.2555
	lds-first			0.5913	0.0342			0.5342	0.0225

For each annotation set, the number of annotations (ref. count) in the test set is given, counting both occurrences (occ.) and unique labels (unique). A dash for precision and recall means that the corresponding system did not predict any unseen concept at all (neither true nor false positive)

difficult sets like PR[_EXT] have comparatively high precision and recall for unseen concepts. A possible explanation is that the former annotation sets have little variability and a high overlap between training and test set, leading to a strong bias for known concepts (overfitting tendency), which is beneficiary for the test set as a whole, but not for the subset of unseen concepts. The latter annotation sets show great variability of concept labels and surface names in the training data, which makes the task harder but also leads to better generalisation, as the classifier cannot achieve good performance by only learning a few frequent concepts.

Error analysis

We performed an analysis of prediction errors in order to find potential weaknesses or systematic mistakes. As expected, many errors are false negatives due to missing training examples. There are several cases where spelled-out mentions are matched, whereas their abbreviated versions are missed. For example, “olfactory tubercle” is correctly linked thanks to the dictionary-based predictions, while the ad-hoc acronym “OT” is missed. False positive predictions are also frequently seen among abbreviations, which have an increased likelihood of being ambiguous. For example, the short-hand “NF” denotes either “neurofilament” or “nuclear factor” in the training set, which cannot always be correctly distinguished by the classifier.

At first sight, it seems like abbreviation expansion should be able to alleviate errors like these. Replacing short forms with their corresponding long forms increases chances for a dictionary match and, since it is performed within document scope, potentially reduces ambiguity. However, abbreviation expansion is not guaranteed to work perfectly and can be a source of confusion even if it does. For example, “OT” was correctly expanded to “olfactory tubercle”. Unfortunately, this misguided the classifier to label the term as *olfactory bulb*, as the first token was only used for this concept in the training data. In our experiments, the net effect of abbreviation expansion was negative, as stated above in the *Hyperparameter tuning* section.

Sometimes, spurious predictions are caused by a substring shared with a training example. Since the WordPiece tokeniser used in (Bio)BERT cuts unknown words into sub-word segments, the classifier sometimes associates a concept label with the fraction of a word, which might trigger false positives in unexpected contexts. As an extreme example, mentions of “PDGFR”, “PFK”, “PKD”, “PI3K”, and “PFKD” are erroneously linked to *phosphoglycerate kinase* (abbreviated “PGK”). This is most likely due to the shared initial letter, as the terms do not refer to semantically similar concepts (even though PFK and PI3K are also kinases). Similarly, “forkhead” is linked to *fork*, “polymorphonuclear” is linked to *nucleus* and “prosensory” is linked to *forebrain* (after the synonym “prosencephalon” seen in training data).

In some cases, the chosen harmonisation strategy prefers an erroneous label over a correct one. For example, the term “monkey” is linked to *mouse* by the ID tagger due to context (training: “mouse kidney”, test: “monkey kidney”). Since the NCBITaxon systems are harmonised with the *ids-first* strategy, this erroneous prediction overrides the correct annotation from the dictionary-based tagger. Conversely, the dictionary predictions for “insulin” always link to PR:000009054, a specific protein. In the ground truth, however, the more general concept PR:000045358 is used throughout the corpus, which denotes a family of proteins. Even though the ID tagger produces correct labels, the *spans-first* strategy used for PR gives precedence to the dictionary predictions in these cases.

Another interesting category of errors are the ones that were amended through the system improvements, i. e. spurious and missing annotations from the shared-task system that are correctly predicted by the current system. A frequent case are short spans by the shared-task system, such as “Ephrin” instead of “Ephrin-B1” for PR or “X” instead of “X-Gal” for CHEBI, which are now correctly recognised. Another re-occurring pattern are incorrect IDs, such as “benzodiazepine” linked to CHEBI:16150 (benzoate) rather than CHEBI:22720 (correct ID by the current system). Furthermore, coverage of frequent terms has improved, for example the shared-task system found “Staphylococcus Aureus” in some context but missed it in others which were correctly identified by the current system.

Conclusions

In this work, we present a concept-recognition architecture for parallel NER and NEN. Compared to a sequential NER+NEN pipeline, our approach avoids error propagation from the span-detection to the normalisation step. Modeling NEN as a sequence-labeling task allows it to operate directly on running text, at the cost of restricting the label

set of the normaliser to the concepts from the training set. We counter these limitations by fusing its predictions with the output of a span detector and a knowledge-based concept recogniser.

In the CRAFT shared task and in the current study, we have shown that parallel concept recognition can outperform a pipeline system created specifically for the CRAFT corpus. Merging the predictions of a span and an ID tagger is a fruitful way of combining the complementary strengths of both of them. However, the specifics of interpolating between span and ID predictions is subject to further research. We took an empiric approach to pick the best harmonisation strategy for each annotation set.

For future work, we intend to test our approach on other datasets. Even though the CRAFT corpus allows validating systems on a broad range of entity types, there is only little opportunity for direct comparison to competing approaches at the time of writing – to the best of our knowledge, there are no published results for the latest version (Version 4) of CRAFT besides the shared task.

Abbreviations

BERT: Bidirectional Encoder Representations from Transformers (neural network); BiLSTM: Bidirectional Long Short-Term Memory (neural network); CHEBI: Chemical Entities of Biological Interest; CL: Cell Ontology; CRAFT: Colorado Richly Annotated Full-Text corpus; CRF: Conditional Random Fields (statistical model); GO_BP: Gene Ontology, biological processes; GO_CC: Gene Ontology, cellular components; GO_MF: Gene Ontology, molecular functions; ID: Identifier; MOP: Molecular Process Ontology; NCBITaxon: US National Center for Biotechnology Information Taxonomy; NEN: Named Entity Normalisation; NER: Named Entity Recognition; PR: Protein Ontology; SER: Slot Error Rate; SO: Sequence Ontology.

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12859-021-04511-y>.

Additional file 1. Hyperparameter tuning. Full results of the hyperparameter-tuning process, performed over the training set in 6-fold cross-validation. The file is an \emph{Open Document Format} table (.ods), which can be viewed and analysed with spreadsheet applications like MS Excel or LibreOffice Calc.

Additional file 2. Harmonisation effect. Tabular version of the values presented in Fig. 5. The file is also in .ods format.

Acknowledgements

We would like to thank the organisers of the CRAFT shared task 2019 for the well-organised competition with high-quality annotations and prompt support.

About this supplement information

This article has been published as part of BMC Bioinformatics Volume 22, Supplement 1 2021: Recent Progresses with BioNLP Open Shared Tasks—Part 2. The full contents of the supplement are available at <https://bmcbioinformatics.biomedcentral.com/articles/supplements/volume-22-supplement-1>.

Authors' contributions

LF conducted the experiments and was a major contributor in writing the manuscript. JC implemented the BioBERT-based system. FR supervised the work and provided guidance on experiments. All authors read and approved the final manuscript.

Funding

This work was supported by the Swiss National Science Foundation [CR3011 162758]; and the Swiss Innovation Agency (InnoSuisse) [25587.2 PFES-ES].

Availability of data and materials

The code and configuration files created for conducting the experiments in the current study are hosted on GitHub, <https://github.com/OntoGene/craft-st>. The trained models for the final results are available from Zenodo, <https://zenodo.org/record/3822363>.

Declarations

Ethics approval and consent to participate

Not applicable

Consent for publication

Not applicable

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Computational Linguistics, University of Zurich, Zurich, Switzerland. ²Dalle Molle Institute for Artificial Intelligence Research (IDSIA USI/SUPSI), Lugano, Switzerland. ³Department of Quantitative Biomedicine, University of Zurich, Zurich, Switzerland. ⁴Swiss Institute of Bioinformatics, Zurich, Switzerland. ⁵Fondazione Bruno Kessler, Trento, Italy.

Received: 29 November 2021 Accepted: 1 December 2021

Published: 24 March 2022

References

- Campos D, Matos S, Oliveira JL. Gimli: open source and high-performance biomedical name recognition. *BMC Bioinformatics*. 2013;14(1):54. <https://doi.org/10.1186/1471-2105-14-54>.
- Xu D, Zhang M, Xie Y, Wang F, Chen M, Zhu KQ, Wei J. DTMiner: identification of potential disease targets through biomedical literature mining. *Bioinformatics*. 2016. <https://doi.org/10.1093/bioinformatics/btw503>.
- Weber L, Münchmeyer J, Rocktäschel T, Habibi M, Leser U. HUNER: improving biomedical NER with pretraining. *Bioinformatics*. 2019;36(1):295–302. <https://doi.org/10.1093/bioinformatics/btz528>.
- Giorgi JM, Bader GD. Towards reliable named entity recognition in the biomedical domain. *Bioinformatics*. 2019;36(1):280–6. <https://doi.org/10.1093/bioinformatics/btz504>.
- Hong SK, Lee J-G. DTranNER: biomedical named entity recognition with deep learning-based label-label transition model. *BMC Bioinformatics*. 2020;21:53. <https://doi.org/10.1186/s12859-020-3393-1>.
- Beltagy I, Lo K, Cohan A. SciBERT: a pretrained language model for scientific text. In *Proceedings of EMNLP-IJCNLP*, 2019; p. 3615–20. <https://doi.org/10.18653/v1/D19-1371>.
- Kim D, Lee J, So CH, Jeon H, Jeong M, Choi Y, Yoon W, Sung M, Kang J. A neural named entity recognition and multi-type normalization tool for biomedical text mining. *IEEE Access*. 2019;7:73729–40. <https://doi.org/10.1109/ACCESS.2019.2920708>.
- Yu X, Hu W, Lu S, Sun X, Yuan Z. BioBERT based named entity recognition in electronic medical record. In: *Proceedings of the 10th international conference on information technology in medicine and education (ITME)*, 2019; p. 49–52. <https://doi.org/10.1109/ITME.2019.00022>.
- Campos D, Matos S, Oliveira JL. A modular framework for biomedical concept recognition. *BMC Bioinformatics*. 2013;14:281. <https://doi.org/10.1186/1471-2105-14-281>.
- D'Souza J, Ng V. Sieve-based entity linking for the biomedical domain. In: *Proceedings of the 53rd annual meeting of the association for computational linguistics and the 7th international joint conference on natural language processing (Volume 2: Short Papers)*, 2015; p. 297–302. <https://doi.org/10.3115/v1/P15-2049>.
- Leaman R, Islamaj Doğan R, Lu Z. DNORM: disease name normalization with pairwise learning to rank. *Bioinformatics*. 2013;29(22):2909–17. <https://doi.org/10.1093/bioinformatics/btt474>.
- Karadeniz İ, Özgür A. Linking entities through an ontology using word embeddings and syntactic re-ranking. *BMC Bioinformatics*. 2019;20(1):156. <https://doi.org/10.1186/s12859-019-2678-8>.
- Hailu ND, Bada M, Hadgu AT, Hunter LE. Biomedical concept recognition using deep neural sequence models. *bioRxiv* 2019. <https://doi.org/10.1101/530337>.
- Pletscher-Frankild S, Jensen LJ. Design, implementation, and operation of a rapid, robust named entity recognition web service. *J Cheminformatics*. 2019;11(1):19. <https://doi.org/10.1186/s13321-019-0344-9>.
- Tseytlin E, Mitchell K, Legowski E, Corrigan J, Chavan G, Jacobson RS. NOBLE—flexible concept recognition for large-scale biomedical natural language processing. *BMC Bioinformatics*. 2016;17(1):1–15. <https://doi.org/10.1186/s12859-015-0871-y>.
- Furrer L, Jancso A, Colic N, Rinaldi F. OGER++: hybrid multi-type entity recognition. *J Cheminformatics*. 2019;11(1):7. <https://doi.org/10.1186/s13321-018-0326-3>.
- Cuzzola J, Jovanović J, Bagheri E. RysannMD: a biomedical semantic annotator balancing speed and accuracy. *J Biomed Inform*. 2017;71:91–109. <https://doi.org/10.1016/j.jbi.2017.05.016>.
- Leaman R, Lu Z. TaggerOne: joint named entity recognition and normalization with semi-Markov Models. *Bioinformatics*. 2016;32(18):2839. <https://doi.org/10.1093/bioinformatics/btw343>.
- ter Horst H, Hartung M, Cimiano P. Joint entity recognition and linking in technical domains using undirected probabilistic graphical models, vol. 10318. Cham: Springer; 2017. p. 166–80. https://doi.org/10.1007/978-3-319-59888-8_15.
- Lou Y, Zhang Y, Qian T, Li F, Xiong S, Ji D. A transition-based joint model for disease named entity recognition and normalization. *Bioinformatics*. 2017;33(15):2363–71. <https://doi.org/10.1093/bioinformatics/btx172>.
- Zhao S, Liu T, Zhao S, Wang F. A neural multi-task learning framework to jointly model medical named entity recognition and normalization. In: *Proceedings of the thirty-third AAAI conference on artificial intelligence (AAAI-19)*, 2019; p. 817–24. <https://doi.org/10.1609/aaai.v33i01.3301817>.
- Bada M, Eckert M, Evans D, Garcia K, Shipley K, Sitnikov D, Baumgartner WA, Cohen KB, Verspoor K, Blake JA, Hunter LE. Concept annotation in the CRAFT corpus. *BMC Bioinformatics*. 2012;13(1):1–20. <https://doi.org/10.1186/1471-2105-13-161>.

23. Cohen KB, Verspoor K, Fort K, Funk C, Bada M, Palmer M, Hunter LE. The Colorado Richly Annotated Full Text (CRAFT) Corpus: Multi-Model Annotation in the Biomedical Domain. Dordrecht: Springer; 2017. p. 1379–94. https://doi.org/10.1007/978-94-024-0881-2_53.
24. Degtyarenko K, de Matos P, Ennis M, Hastings J, Zbinden M, McNaught A, Alcántara R, Darsow M, Guedj M, Ashburner M. ChEBI: a database and ontology for chemical entities of biological interest. *Nucleic Acids Res.* 2008;36(suppl-1):344–50. <https://doi.org/10.1093/nar/gkm791>.
25. Bard J, Rhee SY, Ashburner M. An ontology for cell types. *Genome Biol.* 2005;6(2):21. <https://doi.org/10.1186/gb-2005-6-2-r21>.
26. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, Harris MA, Hill DP, Issel-Tarver L, Kasarskis A, Lewis S, Matese JC, Richardson JE, Ringwald M, Rubin GM, Sherlock G. Gene ontology: tool for the unification of biology. *Nat Genet.* 2000;25(1):25–9. <https://doi.org/10.1038/75556>.
27. Molecular Process Ontology. Processes at the molecular level. <http://obofoundry.org/ontology/mop.html>. Accessed 13 Sep 2021.
28. Federhen S. The NCBI Taxonomy database. *Nucleic Acids Res.* 2012;40(D1):136–43. <https://doi.org/10.1093/nar/gkr1178>.
29. Natale DA, Arighi CN, Barker WC, Blake JA, Bult CJ, Caudy M, Drabkin HJ, D'Eustachio P, Evsikov AV, Huang H, Nchoutmboube J, Roberts NV, Smith B, Zhang J, Wu CH. The Protein Ontology: a structured representation of protein forms and complexes. *Nucleic Acids Res.* 2011;39(suppl-1):539–45. <https://doi.org/10.1093/nar/gkq907>.
30. Eilbeck K, Lewis SE, Mungall CJ, Yandell M, Stein L, Durbin R, Ashburner M. The sequence ontology: a tool for the unification of genome annotations. *Genome Biol.* 2005;6(5):44. <https://doi.org/10.1186/gb-2005-6-5-r44>.
31. Mungall CJ, Torniai C, Gkoutos GV, Lewis SE, Haendel MA. Uberon, an integrative multi-species anatomy ontology. *Genome Biol.* 2012;13(1):5. <https://doi.org/10.1186/gb-2012-13-1-r5>.
32. Baumgartner W, Bada M, Pyysalo S, Ciosici MR, Hailu N, Pielke-Lombardo H, Regan M, Hunter L: CRAFT shared tasks 2019 overview—integrated structure, semantics, and coreference. In: Proceedings of the 5th workshop on BioNLP open shared tasks, 2019; p. 174–84. <https://doi.org/10.18653/v1/D19-5725>.
33. Verspoor K, Cohen KB, Lanfranchi A, Warner C, Johnson HL, Roeder C, Choi JD, Funk C, Malenkiy Y, Eckert M, Xue N, Baumgartner WA, Bada M, Palmer M, Hunter LE. A corpus of full-text journal articles is a robust evaluation tool for revealing differences in performance of biomedical natural language processing tools. *BMC Bioinformatics.* 2012;13(1):207. <https://doi.org/10.1186/1471-2105-13-207>.
34. Crichton G, Pyysalo S, Chiu B, Korhonen A. A neural network multi-task learning approach to biomedical named entity recognition. *BMC Bioinformatics.* 2017;18(1):368. <https://doi.org/10.1186/s12859-017-1776-8>.
35. Groza T, Verspoor K. Assessing the impact of case sensitivity and term information gain on biomedical concept recognition. *PLoS ONE.* 2015;10(3):0119091. <https://doi.org/10.1371/journal.pone.0119091>.
36. Hailu ND. Investigation of traditional and deep neural sequence models for biomedical concept recognition. PhD thesis, University of Colorado at Denver, Anschutz Medical Campus (2019).
37. Basaldella M, Furrer L, Tasso C, Rinaldi F. Entity recognition in the biomedical domain using a hybrid approach. *J Biomed Semant.* 2017;8(1):51. <https://doi.org/10.1186/s13326-017-0157-6>.
38. Yang C-J, Chiang J-H. Gene ontology concept recognition using named concept: understanding the various presentations of the gene functions in biomedical literature. *Database* 2018; <https://doi.org/10.1093/database/bay115>.
39. standoff2conll. Conversion from brat-flavored standoff to CoNLL format. <https://github.com/spyysalo/standoff2conll>. Accessed 3 July 2020.
40. standoff2conll. Forked from spyysalo/standoff2conll. <https://github.com/lfurrer/standoff2conll>. Accessed 3 July 2020.
41. CRAFT shared task evaluation. Code and scripts used for evaluation of the CRAFT Shared Tasks 2019. <https://github.com/UCDenver-ccp/craft-shared-tasks>. Accessed 3 July 2020.
42. Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, Kang J. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics.* 2019. <https://doi.org/10.1093/bioinformatics/btz682>.
43. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of NAACL-HLT 2019, 2019; p. 4171–86. <https://doi.org/10.18653/v1/N19-1423>.
44. Furrer L, Rinaldi F. OGER: OntoGene's entity recogniser in the BeCalm TIPS task. In: Proceedings of the BioCreative V.5 Challenge Evaluation Workshop, 2017; p. 175–182.
45. Sohn S, Comeau DC, Kim W, Wilbur WJ. Abbreviation definition identification based on automatic precision estimates. *BMC Bioinformatics.* 2008;9(1):402. <https://doi.org/10.1186/1471-2105-9-402>.
46. Furrer L, Cornelius J, Rinaldi F. UZH@CRAFT-ST: a sequence-labeling approach to concept recognition. In: Proceedings of the 5th workshop on BioNLP open shared tasks, 2019; p. 185–195. <https://doi.org/10.18653/v1/D19-5726>.
47. Makhoul J, Kubala F, Schwartz R, Weischedel R. Performance measures for information extraction. In: Proceedings of DARPA broadcast news workshop, 1999; p. 249–52.
48. Bossy R, Golik W, Ratkovic Z, Bessières P, Nédellec C. BioNLP shared task 2013—an overview of the bacteria biotope task. In: Proceedings of the BioNLP shared task 2013 workshop, 2013; p. 161–9.
49. Wang JZ, Du Z, Payattakool R, Yu PS, Chen C-F. A new method to measure the semantic similarity of GO terms. *Bioinformatics.* 2007;23(10):1274–81. <https://doi.org/10.1093/bioinformatics/btm087>.
50. Funk C, Baumgartner WA, Garcia B, Roeder C, Bada M, Cohen KB, Hunter LE, Verspoor K. Large-scale biomedical concept recognition: an evaluation of current automatic annotators and their parameters. *BMC Bioinformatics.* 2014;15(1):1–29. <https://doi.org/10.1186/1471-2105-15-59>.
51. Boguslav M, Cohen KB, Baumgartner WA Jr, Hunter LE. Improving precision in concept normalization. In: Pacific symposium on biocomputing 2018, 2018; p. 566–577. https://doi.org/10.1142/9789813235533_0052.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.