

How to tackle an emerging topic?

Combining strong and weak labels for Covid news NER

Aleksander Ficek
University of Waterloo
acficek@uwaterloo.ca

Fangyu Liu
University of Cambridge
f1399@cam.ac.uk

Nigel Collier
University of Cambridge
nhc30@cam.ac.uk

Abstract

Being able to train Named Entity Recognition (NER) models for emerging topics is crucial for many real-world applications especially in the medical domain where new topics are continuously evolving out of the scope of existing models and datasets. For a realistic evaluation setup, we introduce a novel COVID-19 news NER dataset (COVIDNEWS-NER) and release 3000 entries of hand annotated strongly labelled sentences and 13000 auto-generated weakly labelled sentences. Besides the dataset, we propose CONTROSTER, a recipe to strategically combine weak and strong labels in improving NER in an emerging topic through transfer learning. We show the effectiveness of CONTROSTER on COVIDNEWS-NER while providing analysis on combining weak and strong labels for training. Our key findings are: (1) Using weak data to formulate an initial backbone before tuning on strong data outperforms methods trained on only strong or weak data. (2) A combination of out-of-domain and in-domain weak label training is crucial and can overcome saturation when being training on weak labels from a single source.¹

1 Introduction

Named Entity Recognition (NER) is an NLP task that involves identifying key entities in text such as person, location, time or organisation. Research around NER has grown rapidly with the adoption of deep learning techniques and has been an integral step to many NLP pipelines (Sun et al., 2018) such as information retrieval, knowledge base completion, and question answering. As NER models have matured to involve deep Transformer (Vaswani et al., 2017) models and achieve greater performance, the demand for more human labelled strong data has followed. This has become a common bottleneck as attaining more strongly labelled data is expensive and time consuming.

¹Dataset and code is available at <https://github.com/aleksficek/covidnews-ner>.

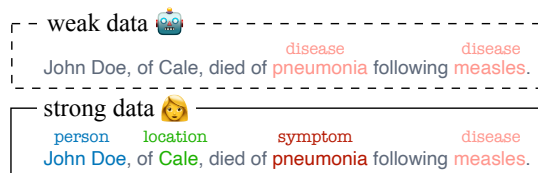


Figure 1: We consider a realistic setup, combining machine-generated noisy weak labels and a small amount of human-generated strong labels for tackling NER in an emerging domain. (Real name replaced with a fictitious name: https://en.wikipedia.org/wiki/John_Doe.)

To work around the limited amount of strongly labelled data, many have experimented with using lower quality weak data generated by weakly supervised methods. Popular techniques to generate weak data include using knowledge bases and heuristic rule based methods while leveraging multiple sources (Lison et al., 2020; Li et al., 2021; Hedderich et al., 2021). All techniques can be applied to any suitable text allowing the methods to generate weak data for any topic. Generating weak labels is especially promising for the medical domain where labelling may require experts to accurately label text and common vocabulary is constantly evolving as seen due to the COVID-19 pandemic. COVID-19 is thus a perfect real-world use case for weakly supervised models.

However, weakly supervised methods are mostly tested on general-domain datasets rather than an emerging domain or topic. The inability of current state-of-the-art NER models to perform when given new biomedical topics such as COVID-19 preprints has been previously stated as a major gap in current NER applications (Langnickel and Fluck, 2021). This was explained by a propensity for models to overfit to currently available training data and a lack of data in the target domain for models to learn such a complex emerging topic like COVID-19. We bridge this gap by proposing a domain-specific NER dataset called COVIDNEWS-NER to evaluate these weakly supervised methods and providing suggestions of combining weak and




Figure 2: CONTROSTER: a double-backboned weak-strong data finetuning architecture.

strong data to address this issue (for a real example of weak and strong labels, see Fig. 1). We expect that the data we publish will drive research around techniques to better adapt to new topics and the strong data we publish to unlock NER applications about COVID-19 and in the medical domain. We plan for our paper to specifically contribute to the relaunched and currently being improved ontology-based text mining tool BioCaster (Meng et al., 2022) for automatic monitoring and surveillance of disease outbreaks.

For our analysis of mixing weak and strong data, we build upon the recent weakly labelled NER model RoSTER (Meng et al., 2021). The model uses a noise robust loss function with noisy label removal, ensemble training and self training without the use of strong data to achieve best performance out of other distantly labelled methods. We propose CONTROSTER, which improves upon RoSTER by performing cross-domain transfer learning over 3 training stages: The first two stages progressively train the model on out-of-domain and in-domain weak data; Afterwards, we finetune the model on in-domain strong data in the last stage (Fig. 2).

To summarise, this paper has the following contributions: (1) We propose a novel COVID-19 NER dataset with 13000 weakly labelled sentences generated by rule based methods and 3000 manually annotated sentences. To our knowledge, this is the first COVID-19 orientated NER dataset in English. (2) We experiment with the data and provide insight into the effect from training with weak and strong data individually and when combined. (3) We propose CONTROSTER, a cross-domain continual training framework, to best leverage strong data and multi-source weak data, and achieve state-of-the-art performance on COVIDNEWS-NER.

2 Dataset: COVIDNEWS-NER

Data generation and filtering. The data consists of 13000 sentences gathered and weakly labelled using the BioCaster ontology-based text mining tool (Collier et al., 2008) with 3000 of the sentences also being manually annotated. BioCaster first generated the text for the dataset by scraping

news articles from multiple local news providers and RSS feeds covering pandemic related topics between approximately January to August 2021. Once BioCaster collected sentences from its news sources, the entries were passed through a text classifier to further refine that selections were oriented towards disease outbreaks (Conway et al., 2009). The classifier selected was recently updated to use the pretrained PubMedBERT (Gu et al., 2021) as a backbone classifier with further finetuning on a binary document classification dataset made of alternating pandemic and normal type news. BioCaster generated part of the dataset from native English texts with additional entries translated from French, Indonesian and Mandarin to English using Language Weaver’s Edge MT engine². The system finally filtered entries by removing entries from duplicate sources from the randomly sampled variety of articles chosen during selection. The weak labels of the dataset were then generated using BioCaster’s rule-based method (Collier et al., 2010). The method is made up of regular expression patterns in simple rule language (SRL), a tool built on top of DIAL (Feldman et al., 2001).

To ensure a high quality of final dataset entries, multiple filtering methods were implemented after this procedure to further prune text with errors. All candidate text was filtered out based on insufficient text lengths, non-ASCII characters involved and text duplicates. Additionally, texts were pruned based on number of grammatical mistakes per entry and finally through manual examination. Manually annotated strong data was labelled by a recent graduate working in the NLP domain. Challenging entries such as differentiating between virus and disease entities were flagged and resolved by discussion with a PhD student and Professor who served as experts in the biomedical NLP domain.³

Dataset entities and content. We employ 10 entity types: Animal, Bacterium, Disease, Location, Organisation, Person, Product, Symptom, Time, Virus.⁴ The Person entity has

²<https://www.rws.com/language-weaver/edge/>

³See Appx. §A.1 for more details.

⁴See Appx. §A.2 for definitions of the entity types.

been expanded to label human cases of a disease as a group of people and the Product entity refers to manufactured articles in the medical domain used during the COVID-19 pandemic (eg. vaccines, face masks). The Virus entity is an especially useful common emerging label since the neighbouring text referencing COVID-19 changes whether it is a Virus or Disease, a common mislabel in weak data. In addition to this, the context developed as a result of the COVID-19 pandemic has produced emerging entities that current rule or knowledge-based labelling systems do not capture. These are however clear to the general public and human annotators which we demonstrate through examples found in COVIDNEWS-NER in Fig. 5. Because the dataset has a heavy focus on the COVID-19 virus, new terminology is also featured surrounding vaccination, testing, variants, etc. Other viruses that gained exposure to the public due to the pandemic are also included in the dataset such as the Zika virus, MERS coronavirus and influenza virus. In general, emerging entities and unique text make COVIDNEWS-NER tailored to pandemics and the medical domain while still providing some useful entities for general NER applications.

Inter-Annotator Agreement (IAA). To demonstrate the quality of the strong data when compared to the weak data we perform an inter-annotator agreement test. The method of Cohen’s Kappa for inter-annotator agreement has been considered inaccurate for NER due to the task not having negative cases to fulfil the methods calculation (Brandesen et al., 2020). In our case, we recruited four additional validators with relevant background to re-annotate 100 randomly selected entries from our dataset. Annotators were given a comprehensive guideline on the labelling strategy and spent on average 90 minutes to read the guideline and complete the labels. We then computed pairwise F1 scores between each of the annotators and the original human annotated 100 strong labels from the dataset. We show this score along with the annotators score when compared to the weak data and the original strong data when compared to the weak data in Tab. 1. The human-labelled strong data has shown high agreement with the validators’ labels, achieving >90% F1 score, demonstrating that the human labelled strong data have high quality.

Data statistics. Tab. 2 summarises multiple metrics that describe the COVIDNEWS-NER dataset. Of the 13000 weak data entries, 3000 of the same

Test	F1	Std. Dev
Strong vs Weak	46.2	-
Weak vs Validators	49.8	2.35
Strong vs Validators	92.3	3.08

Table 1: Dataset Inter-Annotator Agreement

Metric	Weak	Weak-3k	Strong
Total Entries (Sentences)	13000	3000	3000
Total Words	349913	80539	80539
Total Labelled Words	42692	9327	14786
Total Entities	28431	6263	7823
Mean Entity Length	1.50	1.50	1.89
Percent Labelled Words	12.2%	11.6%	18.4%
Mean Entities Per Entry	2.19	2.09	2.61

Table 2: Generic statistics of COVIDNEWS-NER

text have been manually annotated to provide the parallel strong labels with the equivalent 3000 weak labels also evaluated for fair comparison. We provide the total number of words, labelled words and entities for both strong and weak data. There is a noticeable difference in entity length with the average number of words in an entity being 1.489 and 1.890 for the respective 3000 weak and strong data entries. Longer entities are more challenging to fully label and explains how the weak labelling scheme tends to produce shorter labelled entities. This is especially noticeable in types Organisation and Bacterium seen in Tab. 6 where weak data is evaluated directly against strong data. The weak data also had a lower number of entities per entry than the strong data which infers that the weak labelling scheme misses more ambiguous entities and that it is in general under labelled. We provide more analysis on weak vs. strong data in Appx. §A.3 and data split generation in Appx. §A.4.

3 Model

We build upon RoSTER (Meng et al., 2021) which achieves the best performance among distantly-supervised methods. The model contains multiple stages to handle the weak data and its inherent noise. Starting with RoBERTa (Liu et al., 2019) weights, the first step is noise robust training using generalised cross entropy (\mathcal{L}_{GCE}) with tunable parameters dictating noise robustness and noisy label removal. The second step uses ensemble training to improve model stability and the third step introduces contextualised augmentations and self-training with pre-trained RoBERTa embeddings.⁵

⁵See Appx. §B for details of RoSTER.

CONTROSTER: Continually-learned RoSTER. While RoSTER achieves strong performance on noisy data, it remains unclear what is the optimal strategy when both strong and weak data are presented. Additionally, in a real-world use case, we can also assume access to weak labels in other domains. We propose a continual learning approach called CONTROSTER to adapt RoSTER for learning from out of domain weak data and in domain weak and strong data (Fig. 2 represents the fine-tuning pipeline). The pipeline has three training stages: (1) We initially train a RoSTER model on out-of-domain weak data (grey box in Fig. 2). The out-of-domain data are from weak labels generated onto the Wikigold or OntoNotes dataset via knowledge bases (details explained in §4). (2) Then we repeat RoSTER training on in-domain COVIDNews weak data (yellow box) and (3) finally finetuning on strong data with only the noise-robust loss.

4 Experiments

Compared models. We train four models and evaluate them on the COVIDNEWS-NER test set. (1) We train the original RoSTER model with no backbone (initialised with original RoBERTa weights) on strong data. As the data are clean, the model is only trained with the noise robust loss with ensemble learning and self-training stages removed. (2) We train the model on 6000 lines of weak COVIDNews data (this creates a weak COVIDNews backbone) and then finetune on strong COVIDNews data the same as (1). (3) & (4) In the double backbone approach (i.e., our full CONTROSTER model), we first train RoSTER on either the weak labels from the Wikigold dataset (Balasuriya et al., 2009) or the OntoNotes5.0 dataset (Weischedel et al., 2011) followed by training on weak COVIDNews data and finally fine-tuning on strong COVIDNews data. This approach is visualised in Fig. 2 in which (1) is made up of only the Strong COVIDNews box and (2) is made up of the Weak COVIDNews and Strong COVIDNews boxes. The Wikigold dataset contains 13041 lines of training data and 3 overlapping entity categories with COVIDNews out of 4 total entity categories. The OntoNotes5.0 dataset contains 59924 lines of training data and 5 overlapping entity categories with COVIDNews out of 18 total entity categories⁶.

⁶Wikigold/OntoNotes weak data are from Meng et al. (2021).

Main results. Fig. 3 shows the four model’s F1-score performance.⁷ The double backbone approach, i.e. CONTROSTER, performs best for all quantities of strong data used for finetuning. Additionally, using a weak COVIDNews backbone performs noticeably better than the baseline without a backbone for all four amounts of finetuning strong data. The improvement in performance with using either a single or double backbone approach is greatest for experiments with 100 and 500 entries of strong data when compared to using 1000 or 2100. An example of this can be seen by looking at the improvement of 11.0 in F1 score (56.7 to 66.7) when the weak Wiki+COV. backbone was paired with 100 entries of strong data in comparison to an increase of 2.2 in F1 score (74.6 to 76.8) in the same scenario when 2100 entries of strong data were used.




Figure 3: Main weak and strong data results. When using different number of strong data points, having the COVIDNews backbone (COV.) has always helped. Adding the Wiki/Onto backbone on top of COV. has also almost universally helped.

These findings support the use of combining weak data with strong data through the method of transfer learning for research and NER applications. We provide insight into the amounts of strong data necessary for effective combination of the two types of data. We recommend using a weak data generated backbone in general NER models with the potential for profound impact in few-shot learning models that have a limited number of strong data. Similarly, in scenarios where only weak data is available we advise manually annotating a minimum of 100 sentences can lead to large improvements in NER model performance.

⁷Precision & Recall follow the same trend (see Tab. 8).




Figure 4: Weak data performance of RoSTER.

Weak data study. We investigate the magnitude of weak data for the models trained to be effective on the COVIDNEWS-NER dataset. Fig. 4 shows the improvement in precision, recall and F1 scores as weak data scales with the RoSTER model. After surpassing the 6k entry mark, the weak data saturates and does not lead to improved performance for each additional entry. Alternatively, the strong data in Tab. 8 demonstrates how the higher quality strong labels continue to scale with additional data. Varying amounts of weak data were also evaluated for the COVIDNews backbone prior to finetuning on strong data. Tab. 10 in Appendix shows tests with 2000, 6000 and 10000 weak entries paired with 100, 500, 1000 and 2100 strong entries. Although using a weak data backbone was clearly beneficial, the amount of weak data to train the backbone was fairly insignificant.

5 Further Discussion

Combining weak & strong data. We believe a wide and deep strategy should be used when combining weak & strong data. This means training on weak data to embed the model with an expansive breadth of knowledge over all entities and then fine-tuning on strong data to overwrite noise in the weak data and generate more sophisticated ability in NER. We suggest two avenues for further research to maximise utility of weak data when paired with strong data. An improvement in noise reduction techniques via improved loss functions or model architecture will serve beneficial in allowing model performance to continue to scale with increases in weak data. Additionally, investigation into how weak data generated from specific rules saturates individually and after fine-tuning with strong data will also lead to improved knowledge on mixing the different forms of labels.

Out-of-domain weak data. We observed that using the two stage backbone pipeline of first training on a different NER dataset and then following

through with training on COVIDNEWS-NER weak and strong data led to a further improvement in performance across varying COVIDNEWS-NER weak and strong data lengths. Even though the Wikigold dataset had only three overlapping entity categories with COVIDNEWS-NER and OntoNotes5.0, the performance still improved in part due to the overlapping categories being the prominent ones in COVIDNEWS-NER (eg. organisation, person, location, date). The additional variety of labels in the crossover entities led to improved precision, recall and F1 scores in those categories and overall scores. We implore future works to evaluate the impact of overlapping and non-overlapping categories from cross-domain weak data backbones while considering overlapping category definition similarities and differences. Overall, similar to how Li et al. (2021) determined the importance of using multiple sources for text to be distantly labelled, we conclude it is also beneficial to use different weak labelling techniques to create a diverse collection of weak data. We implore future work to investigate the diversity necessary for optimal combination of weak data from different sources.

In Appendix, we include a dedicated related work section (Appx. §D) for interested readers.

6 Conclusion

We presented COVIDNEWS-NER, an English COVID-19 Named Entity Recognition dataset in the pandemic news domain, addressing current NER models’ lack of ability to tackle new and out-of-domain topics. We labelled 13000 entries using a rule-based system to generate weak labels and 3000 entries using hand annotation to generate strong labels. We further proposed a continual learning approach called CONTROSTER that transfers knowledge learned in both out-of-domain and in-domain weak data. After finetuning on strong in-domain data, CONTROSTER achieved state-of-the-art performance on our proposed dataset. We further provide detailed and thorough analysis into how to successfully combine both types of data and suggest promising avenues for future research. We think that the dataset we provide and the findings we conclude will be beneficial to other NER applications, such as improving the evaluation and ability of the BioCaster pandemic surveilling tool. We hope that our work drives more research in leveraging a combination of weak and strong data to improve performance on new topics such as the COVID-19 pandemic.

Acknowledgements

We are grateful to RWS Language Weaver for use of their neural MT engine. We also thank Qianchu Liu, Parth Shah, Chandni Bhatt and Marko Popovic for contributing to the inter-annotator agreement.

References

- Dominic Balasuriya, Nicky Ringland, Joel Nothman, Tara Murphy, and James R. Curran. 2009. [Named entity recognition in Wikipedia](#). In *Proceedings of the 2009 Workshop on The People's Web Meets NLP: Collaboratively Constructed Semantic Resources (People's Web)*, pages 10–18, Suntec, Singapore. Association for Computational Linguistics.
- Alex Brandsen, Suzan Verberne, Milco Wansleeben, and Karsten Lambers. 2020. [Creating a dataset for named entity recognition in the archaeology domain](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4573–4577, Marseille, France. European Language Resources Association.
- Rosario Catelli, Francesco Gargiulo, Valentina Casola, Giuseppe De Pietro, Hamido Fujita, and Massimo Esposito. 2020. [Crosslingual named entity recognition for clinical de-identification applied to a covid-19 italian data set](#). *Applied Soft Computing*, 97:106779.
- Nigel Collier, Son Doan, Ai Kawazoe, Reiko Matsuda Goodwin, Mike Conway, Yoshio Tateno, Quoc-Hung Ngo, Dinh Dien, Asanee Kawtrakul, Koichi Takeuchi, et al. 2008. [Biocaster: detecting public health rumors with a web-based text mining system](#). *Bioinformatics*, 24(24):2940–2941.
- Nigel Collier, Reiko Matsuda Goodwin, John McCrae, Son Doan, Ai Kawazoe, Mike Conway, Asanee Kawtrakul, Koichi Takeuchi, and Dinh Dien. 2010. [An ontology-driven system for detecting global health events](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 215–222, Beijing, China. Coling 2010 Organizing Committee.
- Mike Conway, Son Doan, Ai Kawazoe, and Nigel Collier. 2009. [Classifying disease outbreak reports using n-grams and semantic features](#). *International journal of medical informatics*, 78(12):e47–e58.
- Ronen Feldman, Yonatan Aumann, Yair Liberzon, Kfir Ankori, Jonathan Schler, and Benjamin Rosenfeld. 2001. [A domain independent environment for creating information extraction modules](#). In *Proceedings of the Tenth International Conference on Information and Knowledge Management, CIKM '01*, page 586–588, New York, NY, USA. Association for Computing Machinery.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Domain-specific language model pretraining for biomedical natural language processing](#). *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- Michael A Hedderich, Lukas Lange, and Dietrich Klakow. 2021. [Anea: distant supervision for low-resource named entity recognition](#). *arXiv preprint arXiv:2102.13129*.
- Haoming Jiang, Danqing Zhang, Tianyu Cao, Bing Yin, and Tuo Zhao. 2021. [Named entity recognition with small strongly labeled and large weakly labeled data](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1775–1789, Online. Association for Computational Linguistics.
- Lisa Langnickel and Juliane Fluck. 2021. [We are not ready yet: limitations of transfer learning for disease named entity recognition](#). *bioRxiv*.
- Yinghao Li, Pranav Shetty, Lucas Liu, Chao Zhang, and Le Song. 2021. [BERTifying the hidden Markov model for multi-source weakly supervised named entity recognition](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6178–6190, Online. Association for Computational Linguistics.
- 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*, pages 1054–1064.
- Pierre Lison, Jeremy Barnes, Aliaksandr Hubin, and Samia Touileb. 2020. [Named entity recognition without labelled data: A weak supervision approach](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1518–1533, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Yu Meng, Yunyi Zhang, Jiabin 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](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10367–10378, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Zaiqiao Meng, Anya Okhmatovskaia, Maxime Polleri, Yannan Shen, Guido Powell, Zihao Fu, Iris Ganser, Meiru Zhang, Nicholas B King, David Buckeridge, et al. 2022. *Biocaster in 2021: automatic disease outbreaks detection from global news media*. *Bioinformatics*.

Peng Sun, Xuezheng Yang, Xiaobing Zhao, and Zhijuan Wang. 2018. *An overview of named entity recognition*. In *2018 International Conference on Asian Language Processing (IALP)*, pages 273–278. IEEE.

Thinh Hung Truong, Mai Hoang Dao, and Dat Quoc Nguyen. 2021. *COVID-19 named entity recognition for Vietnamese*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2146–2153, Online. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. *Attention is all you need*. *Advances in neural information processing systems*, 30.

Ralph Weischedel, Sameer Pradhan, Lance Ramshaw, Martha Palmer, Nianwen Xue, Mitchell Marcus, Ann Taylor, Craig Greenberg, Eduard Hovy, Robert Belvin, et al. 2011. *Ontonotes release 4.0*. *LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium*.

Shi Zong, Ashutosh Baheti, Wei Xu, and Alan Ritter. 2020. *Extracting a knowledge base of covid-19 events from social media*. *arXiv preprint arXiv:2006.02567*.

A More Dataset Details

A.1 Data Generation and Filtering Details

The text used in the dataset and the corresponding manually annotated strong labels had some deviation in metrics depending on the original language translated from. Additional statistics of the dataset for each text language origin can be seen in [Tab. 3](#). The metrics suggest that Mandarin was an especially useful and unique language to generate data from as it introduced many longer entities due to person titles and location addresses being more verbose. Other languages like French introduced text with fewer entities per sentence than others.

Metric	Total	Eng.	Fre.	Ind.	Man.
Total Entries	3000	1500	505	500	495
Mean Entity Length	1.89	1.79	1.80	1.68	2.46
Percent Labelled Words	18.4%	17.5%	15.4%	18.9%	22.9%
Mean Entities Per Entry	2.61	2.71	2.09	2.70	2.74

Table 3: Generic statistics of COVIDNEWS-NER strong data separated by language

For filtering, texts that were less than 4 words, less than 15 characters or greater than 500 characters were removed. Duplicate sentences were also filtered out and poorly structured entries were identified using the LanguageTool grammar checking API⁸. This checked and removed entries with grammar, punctuation and syntactical mistakes.

A.2 Entity Definitions

The exact definitions of the 10 entity types included in COVIDNEWS-NER can be found in [Tab. 4](#).

A.3 Weak vs. Strong Data

We further investigate the differences in performance and style between the rule based weak data generation method and human annotated strong data. [Tab. 5](#) shows the difference in entity count for each category between the two types of data. Organisation is noticeably out numbered in the strong data case which can be explained by the category requiring more in depth understanding of contextual knowledge as that can change it being classified as a Location or Organisation (eg. “The White House”). Another notable difference is that larger groups of words are categorised as entities in the strong labels when compared to the weak. The strong data contains 2.5 times more entities containing greater than three words due to the difficulty in labelling longer entities. The Symptom, Disease and Virus entity categories in the dataset significantly orient the dataset towards the COVID-19 pandemic. The difference between the three categories are challenging to distinguish, examples of which are shown in [Fig. 5](#). [Tab. 6](#) shows a detailed breakdown of the performance of weak data when evaluated directly against strong data and [Tab. 7](#) shows the performance across the different languages the text was translated from.

A.4 Data Split

To best split the data into training, validation and test sets, a unique Monte Carlo technique was implemented to insure entities with limited labels such as *Bacterium* were allocated in proper amounts to each partition. An optimal distribution of all entities was generated and 10000 random iterations of the input data was split and scored based on fractional proximity to the optimal distribution. The split dataset partitions had 2100/300/600 entries in train, validation and test sets respectively.

⁸<https://languagetool.org/>

Entity Type	Definition
Animal	Multi-cell organisms that are eukaryotes of the kingdom Animalia, other than humans.
Bacterium	Single-celled prokaryotic microorganisms of the bacteria domain.
Disease	A disorder of a structure or function that affects an organism, associated with specific phenotypes.
Location	A politically or geographically defined location for example a region, a province, a town.
Organisation	Named corporate, governmental, or other organisational entity.
Person	A person or group of persons.
Product	Medical articles or substances manufactured and used throughout pandemics.
Symptom	Phenotypic descriptions of any abnormal morphology, physiology or behaviour.
Time	Temporal expressions that can be anchored on a timeline.
Virus	A disease causing infectious agent that is non-living.

Table 4: Entity Type Definitions

Figure 5 shows three examples of text snippets from the COVIDNEWS-NER dataset, illustrating the difference between weak and strong data. Each example consists of a 'weak data' snippet (dashed border) and a 'strong data' snippet (solid border). The 'strong data' snippets have colored labels above the words they identify: 'organisation' (purple), 'virus' (pink), 'time' (blue), 'disease' (red), and 'symptom' (orange).

- Example 1:** The CDC points out that this also shows the power of Delta's mutation. (Strong data labels: organisation, virus)
- Example 2:** Early Monday morning, Forber, 57, succumbed to pneumonia caused by COVID-19. (Strong data labels: time, disease, disease)
- Example 3:** Among the specimens collected during December 16 - 23, 22 were the B.1.351 variant. (Strong data labels: time, time, virus)

Figure 5: Additional examples from the COVIDNEWS-NER dataset

Entity Type	Strong	Weak
Animal	177	201
Bacterium	25	12
Disease	641	612
Location	1703	1568
Organisation	1076	270
Person	2652	2370
Product	233	203
Symptom	121	146
Time	799	697
Virus	396	184
Total	7823	6263

Table 5: Entity counts in COVIDNEWS-NER

Entity Type	Pre.	Rec.	F1	Support
Animal	62.2	70.6	66.1	177
Bacterium	33.3	16.0	21.6	25
Disease	66.2	63.2	64.6	641
Location	57.0	52.4	54.6	1703
Organisation	33.3	8.4	13.4	1076
Person	46.7	41.7	44.0	2652
Product	63.1	54.9	58.7	233
Symptom	46.6	56.2	50.9	121
Time	68.4	59.7	63.8	799
Virus	49.5	23.0	31.4	396
Weighted Avg	51.8	43.3	46.2	7823

Table 6: Weak Data Performance

Entry Language	Entries	Pre.	Rec.	F1
Combined	3000	59.4	44.3	49.7
English	1500	60.3	41.0	47.5
French	505	62.3	53.3	56.5
Indonesian	500	62.6	49.4	54.1
Mandarin	495	53.9	42.6	46.1

Table 7: Weak Data Performance Across Languages

B RoSTER Details

Here we explain the RoSTER methodology in greater detail. The first stage is known as the noise-robust learning stage and introduces two hyperparameters for adjusting to noisy labels. (Meng et al., 2021) uncovers that cross entropy loss is useful for model convergence but is sensitive to noise while mean absolute error (MAE) loss is robust to noise at the cost of convergence. The generalised cross entropy loss uses a q parameter to adjust cross entropy towards CE loss by lowering and towards MAE loss by raising. A thresholding parameter is introduced to remove incorrect labels during the training process. The parameter τ is used as a threshold for comparing model predictions with distant labels. If there are differences between model predictions and distant labels greater than the threshold, the model omits those labels when updating weights.

$$\mathcal{L}_{GCE} = \sum_{i=1}^n w_i \frac{1 - f_{i,y_i}(x; \theta)^q}{q} \quad (1)$$

RoSTER also implements ensemble and self training stages to improve results on distantly labelled data. The ensemble stage uses a K parameter to determine the number of models trained using different seeds and a final model is employed to approximate the performance of trained models by minimising Kullback–Leibler (KL) divergence loss. Prior to self-training, contextualised augmentations are generated using PLM’s like RoBERTa. Then

the model trains on an unlabelled version of the corpus to leverage knowledge embedded in the selected PLM while generalising model predictions to tokens removed by noisy label removal. Self-training is done by polarising predictions during iterations by squaring high-confidence predictions and normalising low-confidence predictions.

C More Experimental Details

Weak-strong main results full table (Tab. 8). In the main text we showed performance of CON-TROSTER and its ablated versions’ in Fig. 3. Here, we provide a more detailed view of the same data, listing also Precision and Recall scores in Tab. 8.

W. Backbone	S. Tuning	Pre.	Rec.	F1
None	100	50.9	65.2	56.7
COV.	100	59.9	71.7	65.2
Wiki. + COV.	100	62.1	72.5	66.7
Onto. + COV.	100	60.2	69.9	64.6
None	500	62.8	73.0	67.3
COV.	500	66.3	73.9	69.7
Wiki. + COV.	500	68.0	75.7	71.6
Onto. + COV.	500	68.5	76.7	72.2
None	1000	66.7	75.8	70.9
COV.	1000	69.1	76.1	72.3
Wiki. + COV.	1000	69.6	76.6	72.8
Onto. + COV.	1000	70.2	77.3	73.5
None	2100	71.7	77.9	74.6
COV.	2100	72.9	78.2	75.4
Wiki. + COV.	2100	73.9	79.9	76.7
Onto. + COV.	2100	74.2	79.7	76.8

Table 8: Main weak and strong data results

Weak backbone saturation data (Tab. 9). Since we have an in-domain rule-based weak labeller, why not generate as much in-domain weak data as possible? As mentioned in the main text Fig. 4, we found that in-domain weak data only helps up to a certain point. Here we list the exact precision, recall and F1 results used for plotting the figure of reference in the main text (Tab. 9).

Weak Data	Pre.	Rec.	F1
1000 entries	46.4	52.7	47.7
2000 entries	47.1	52.4	48.7
4000 entries	48.9	52.4	49.8
6000 entries	52.2	52.2	51.2
10000 entries	51.8	50.5	50.3
13000 entries	49.7	49.4	48.7

Table 9: Performance of RoSTER when varying number of weak data

Weak data study (§4) full table (Tab. 10). In the main text we discussed varying amounts of weak

data when pretraining on COVIDNEWS-NER. Here we attach the full table (Tab. 10) for reference.

W. Backbone	S. Tuning	Pre.	Rec.	F1
None	100	50.9	65.2	56.7
Weak 2000	100	59.9	69.2	64.1
Weak 6000	100	60.4	70.6	65.0
Weak 10000	100	61.5	70.9	65.7
None	500	62.8	73.0	67.3
Weak 2000	500	67.1	74.2	70.3
Weak 6000	500	66.3	73.9	69.7
Weak 10000	500	66.9	74.3	70.2
None	1000	66.7	75.8	70.9
Weak 2000	1000	69.6	76.1	72.6
Weak 6000	1000	69.1	76.1	72.3
Weak 10000	1000	70.2	76.7	73.2
None	2100	71.7	77.9	74.6
Weak 2000	2100	72.7	77.2	74.7
Weak 6000	2100	72.9	78.2	75.4
Weak 10000	2100	72.6	77.5	74.9

Table 10: Weak data quantities with strong data results

D Related Work

Our work is related to other COVID-19 datasets in the NER domain. [Truong et al. \(2021\)](#) introduced a COVID-19 NER dataset for the low resource language of Vietnamese and [Catelli et al. \(2020\)](#) provided a NER dataset based on medical records in Italian. [Zong et al. \(2020\)](#) scraped and annotated COVID-19 related tweets, generating a knowledge base but labelling events (eg. tested positive, can not test) as opposed to entities necessary for NER. Our dataset provides the first COVID-19 NER dataset in English with distantly supervised weak data and human annotated strong data.

Our analysis of combining weak and strong data is related to previous methods which successfully utilise either types of data to improve performance. [Liang et al. \(2020\)](#) implemented the use of pre-trained language models with subsequent self-training with weak labels generated through knowledge bases to improve model performance. [Jiang et al. \(2021\)](#) architected a multistage pipeline involving pre-training on unlabelled data, weak label completion, a noise robust loss function and fine tuning on strong data to effectively null the impact of noise. We build upon these works and provide insight into using the two forms of data together in addition to using cross-domain datasets on an emerging topic such as COVID-19.