



Models and Processes to Extract Drug-like Molecules From Natural Language Text

Zhi Hong^{1*}, J. Gregory Pauloski¹, Logan Ward², Kyle Chard^{1,2,3}, Ben Blaiszik^{2,3} and Ian Foster^{1,2,3}

¹Department of Computer Science, University of Chicago, Chicago, IL, United States, ²Data Science and Learning Division, Argonne National Laboratory, Lemont, IL, United States, ³Globus, University of Chicago, Chicago, IL, United States

OPEN ACCESS

Edited by:

James Leland Olds,
George Mason University,
United States

Reviewed by:

Valentina Tozzini,
Consiglio Nazionale delle Ricerche,
Italy
Shang Gao,
Oak Ridge National Laboratory (DOE),
United States

*Correspondence:

Zhi Hong
hongzhi@uchicago.edu

Specialty section:

This article was submitted to
Biological Modeling and Simulation,
a section of the journal
Frontiers in Molecular Biosciences

Received: 30 November 2020

Accepted: 11 August 2021

Published: 30 August 2021

Citation:

Hong Z, Pauloski J G, Ward L,
Chard K, Blaiszik B and Foster I (2021)
Models and Processes to Extract
Drug-like Molecules From Natural
Language Text.
Front. Mol. Biosci. 8:636077.
doi: 10.3389/fmolb.2021.636077

Researchers worldwide are seeking to repurpose existing drugs or discover new drugs to counter the disease caused by severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2). A promising source of candidates for such studies is molecules that have been reported in the scientific literature to be drug-like in the context of viral research. However, this literature is too large for human review and features unusual vocabularies for which existing named entity recognition (NER) models are ineffective. We report here on a project that leverages both human and artificial intelligence to detect references to such molecules in free text. We present 1) a iterative model-in-the-loop method that makes judicious use of scarce human expertise in generating training data for a NER model, and 2) the application and evaluation of this method to the problem of identifying drug-like molecules in the COVID-19 Open Research Dataset Challenge (CORD-19) corpus of 198,875 papers. We show that by repeatedly presenting human labelers only with samples for which an evolving NER model is uncertain, our human-machine hybrid pipeline requires only modest amounts of non-expert human labeling time (tens of hours to label 1778 samples) to generate an NER model with an F-1 score of 80.5%—on par with that of non-expert humans—and when applied to CORD-19, identifies 10,912 putative drug-like molecules. This enriched the computational screening team's targets by 3,591 molecules, of which 18 ranked in the top 0.1% of all 6.6 million molecules screened for docking against the 3CLPro protein.

Keywords: coronavirus disease-19, coronavirus disease-19 open research dataset challenge, named entity recognition, long-short term memory, data mining

1 INTRODUCTION

The Coronavirus Disease (COVID-19) pandemic, caused by transmissible infection of the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), has resulted in tens of millions of diagnosed cases and over 1,450,000 deaths worldwide (Dong et al., 2020); straining healthcare systems, and disrupting key aspects of society and the wider economy. It is thus important to identify effective treatments rapidly via discovery of new drugs and repurposing of existing drugs. Here, we leverage advances in natural language processing to enable automatic identification of drug candidates being studied in the scientific literature.

The magnitude of the pandemic has resulted in an enormous number of academic publications related to COVID-19 research since early 2020. Many of these articles are collated in the COVID-19

Open Research Dataset Challenge (CORD-19) collection (Allen Institute For AI, 2020; Wang et al., 2020). With 198,875 articles at the time of writing, that collection is far too large for humans to read. Thus, tools are needed to automate the process of extracting relevant data, such as drug names, testing protocols, and protein targets. Such tools can save domain experts significant time and effort.

Extracting named entities from scientific texts has been studied for more than 2 decades. Prior work relied on matching tokens (words) in text to entries in existing databases or ontologies (Rindflesch et al., 1999; Furrer et al., 2019). However, for our task, existing drug databases like DrugBank cover both *too much*, in that they include entities of many types (for example, “rabbit” is in DrugBank as an allergen), and *too little*: because the creation of such databases is time-consuming, they cannot keep up with the new entities in the latest COVID-19-related publications (~200 k papers in the CORD-19 corpus), which is exactly what we want to extract here.

Machine learning (ML) and deep learning (DL) methods have also been used for identifying named entities from free text. While less dependent on a comprehensive ontology, they need labeled data for training. Traditional training data collection employs a brute-force approach, collecting a large corpus and labeling each and every word indiscriminately (Bada et al., 2012). The time and human resources involved is enormous. We, in contrast, had access to just a few assistants who worked on labeling only during free time in their day job. Thus we needed a more careful approach of selective labeling to maximize their effectiveness.

Towards this goal, we describe here how we have tackled two important problems: creating labelled training data via judicious use of scarce human expertise, and applying a named entity recognition (NER) model to automatically identify drug-like molecules in text. We are looking not only for novel drugs under development, but also any small molecule drug that has been used to treat patients with COVID-19 or similar infectious diseases like SARS and MERS. In the absence of expert-labeled data for the growing COVID literature, we employ an iterative model-in-the-loop collection process inspired by our previous work (Tchoua et al., 2019a,b) and demonstrate that it can build a high quality training set without input from domain scientists. We first assemble a small bootstrap set of human-verified examples to train a model for identifying similar examples. We then iteratively apply the model, use human reviewers to verify the predictions for which the model is least confident, and retrain the model until the improvement in performance is less than a threshold. (The human reviewers were administrative staff without scientific backgrounds, with time available for this task due to the pandemic.)

Having collected adequate training data via this model-guided human annotation process, we then use the resulting labeled data to re-train a NER model originally developed to identify polymer names in materials science publications (Hong et al., 2020b) and apply this trained model to CORD-19. We show that the labeled data produced by our approach are of sufficiently high quality than when used to train NER models, which achieves a best F-1 score of 80.5%—roughly equivalent to that achieved by non-expert humans.

The labeled data, model, and model results are all available online, as described in **Section 5**.

2 MATERIALS AND METHODS

We aim to develop and apply new computational methods to mine the scientific literature to identify small molecules that have been investigated or found useful as antiviral therapeutics. For example, processing the following sentence should allow us to determine that the drug *sofosbuvir* has been found effective against the Zika virus: “Sofosbuvir, an FDA-approved nucleotide polymerase inhibitor, can efficiently inhibit replication and infection of several ZIKV strains, including African and American isolates.” (Bullard-Feibelman et al., 2017).

This problem of identifying drug-like molecules in text can be divided into two linked problems: 1) identifying references to small therapeutic molecules (“drugs”) and 2) determining what the text says about those molecules. In this work, we consider potential solutions to the first problem.

A simple way to identify entities in text that belong to a specialized class (e.g., drug-like molecules) is to refer to a curated list of valid names, if such is available. In the case of drugs, we might think to use DrugBank (Wishart et al., 2018) or the FDA Drug Database (Center for Drug Evaluation and Research, 2020), both of which in fact list *sofosbuvir*. However, such databases are not in themselves an adequate solution to our problem, for at least two reasons. First, they are rarely complete. The tens of thousands of entity names in DrugBank and the FDA Drug Database together are just a tiny fraction of the billions of molecules that could potentially be used as drugs. Second, such databases may be overly general: DrugBank, for example, includes the terms “rabbit” and “calcium,” neither of which have value as antiviral therapeutics. In general, the use of any such list to identify entities will lead to both false negatives and false positives. We need instead to employ the approach that a human reader might follow in this situation, namely to scan text for words that appear in contexts in which a drug name is likely to appear. In the following, we explain how we combine human and artificial intelligence for this purpose.

2.1 Automated Drug Entity Extraction From Literature

Finding strings in text that refer to drug-like molecules is an example of Named Entity Recognition (NER) (Nadeau and Sekine, 2007), an important NLP task. Both grammatical and statistical (e.g., neural network-based) methods have been applied to NER; the former can be more accurate, but require much effort from trained linguists to develop. Statistical methods use supervised training on labeled examples to learn the contexts in which entities of interest (e.g., drug-like molecules) are likely to occur, and then classify previously unseen words as such entities if they appear in similar contexts. For instance, a training set may contain the sentence “Ribavirin was administered once daily by the i. p. route” (Oestereich et al., 2014), with *ribavirin* labelled as

Drug. With sufficient training data, the model may learn to assign the label *Drug* to *arbidol* in the sentence “Arbidol was administered once daily per os using a stomach probe” (Oestereich et al., 2014). This learning approach can lead to general models capable of finding previously unseen candidate molecules in natural language text.

The development of effective statistical NER models is complicated by the many contexts in which names can occur. For example, while the contexts just given for *ribavirin* and *arbidol* are similar, both are quite different from that quoted for *sofosbuvir* earlier. Furthermore, authors may use different wordings and sentence structures: e.g., “given by i. p. injection once daily” rather than “administered once daily by the i. p. route.” Thus, statistical NER methods need to do more than learn template word sequences: they need to learn more abstract representations of the context(s) in which words appear. Modern NLP and NER systems do just that (Chiu and Nichols, 2016).

2.1.1 SpaCy and Keras-Long-Short Term Memory Models

We consider two NER models in this paper, SpaCy and a Keras long-short term memory (LSTM) model. Both models are publicly available on DLHub (Li et al., 2021) and GitHub, as described in Section 5.

SpaCy (Honnibal and Montani, 2020a; Honnibal et al., 2020) is an open source NLP library that provides a pre-trained entity recognizer that can recognize 18 types of entities, including PERSON, ORGANIZATION, LOCATION, and PRODUCT. Its model calculates a probability distribution of a word over the entity types, and outputs the type with the highest probability as the predicted type for that word. When pre-trained on the OntoNotes five dataset of over 1.5 million labeled words (Weischedel et al., 2013), the SpaCy entity recognizer can identify supported entities with 85.85% accuracy. However, it does not include drug names as a supported entity class, and thus we would need to retrain the SpaCy model on a drug-specific training corpus. Unfortunately, there is no publicly available corpus of labeled text for drug-like molecules in context. Thus, we need to use other methods to retrain this model (or other NER models), as we describe in Section 4.

While SpaCy is easy to use, it lacks flexibility: its end-to-end encapsulation does not expose many tunable parameters. Thus we also explore the use of a Keras-LSTM model that we developed in previous work for identification of polymers in materials science literature (Hong et al., 2020b). This model is based on the Bidirectional LSTM network with a conditional random field (CRF) layer added on top. It takes training data labeled according to the “IOB” schema. The first word in an entity is given the label “B” (Beginning), the following words in the same entity are labeled “I” (Inside), and non-entity words are labeled “O” (outside). During prediction, the Bi-LSTM network tries to assign one of “IOB” to each word in the input sentence, but it has no awareness of the validity of the label sequence. The CRF layer is used on top of Bi-LSTM to lower the probability of invalid label sequences (e.g., “OIO”).

We compare the performance of SpaCy and Keras-LSTM models under various conditions in Section 2.2.

2.1.2 Model-In-The-Loop Annotation Workflow

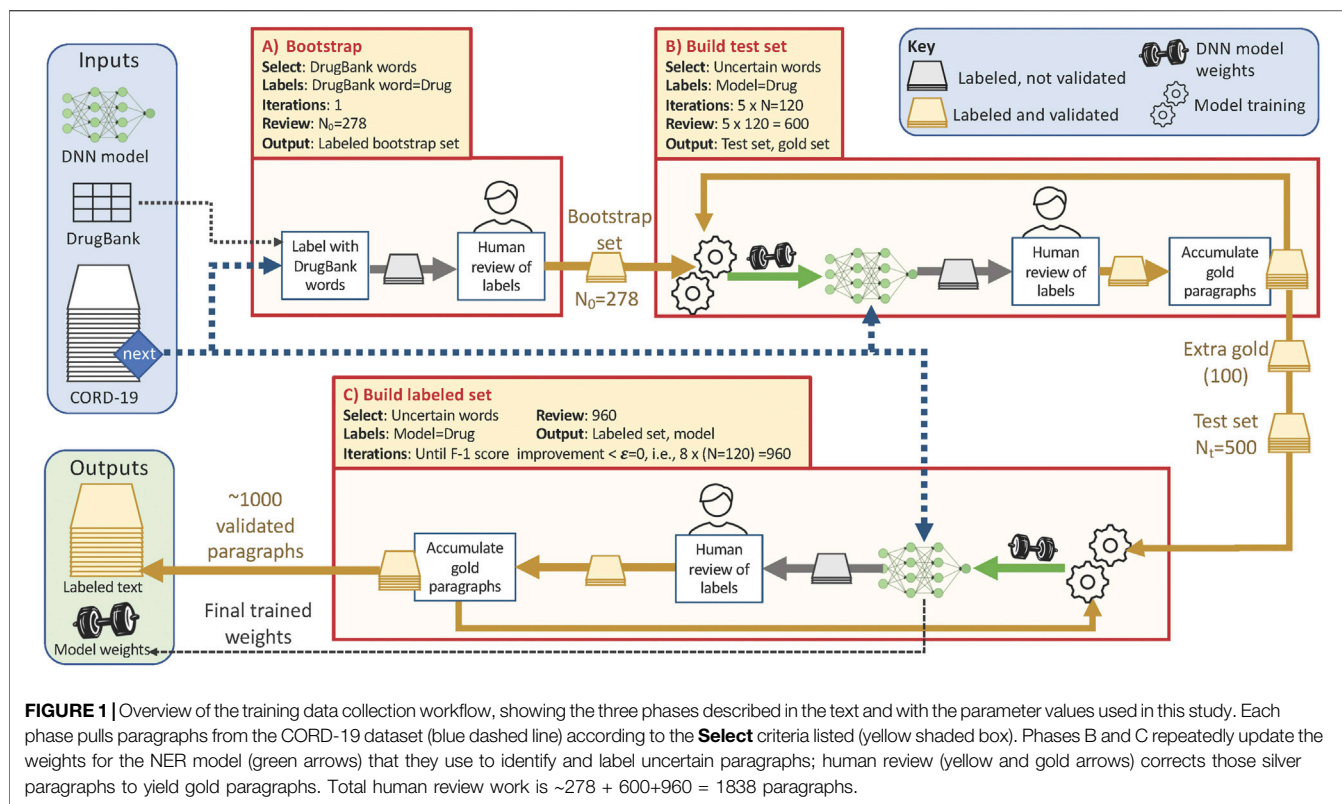
We address the lack of labeled training data by using Algorithm 1 (and see Figure 1) to assemble a set of human- and machine-labeled data from CORD-19 (Wang et al., 2020). In describing this process, we refer to paragraphs labeled automatically via a heuristic or model as *silver* and to silver paragraphs for which labels have been corrected by human reviewers as *gold*. We use the Prodigy machine learning annotation tool to manage the review process: reviewers are presented with a silver paragraph, with putative drug entities highlighted; they click on false negative and false positive words to add or remove the highlights and thus produce a gold paragraph. Prodigy saves the corrected labels in standard NER training data format.

Our algorithm involves three main phases, as follows. In the first **bootstrap** phase, we assemble an initial test set of gold paragraphs for use in subsequent data acquisition. We create a first set of silver paragraphs by using a simple heuristic: we select N_0 paragraphs from CORD-19 that contain one or more words in DrugBank with an Anatomical Therapeutic Chemical Classification System (ATC) code, label those words as drugs, and ask human reviewers to correct both false positives and false negatives in our silver paragraphs, creating gold paragraphs. In the subsequent **build test set** phase, we repeatedly use all gold paragraphs obtained so far to train an NER model; use that model to identify and label additional silver paragraphs, and engage human reviewers to correct false positives and false negatives, creating additional gold paragraphs. We repeat this process until we have N_t initial gold paragraphs.

In the third **build labeled set** phase, we repeatedly use an NER model trained on all human-validated labels obtained to date, with the N_t gold paragraphs from the bootstrap phase used as a test set, to identify and label promising paragraphs in CORD-19 for additional human review. To maximize the utility of this human effort, we present the reviewers only with paragraphs that contain one or more *uncertain words*, i.e., words that the NER model identifies as drug/non-drug with a confidence in the range (min, max). We continue this process of model retraining, paragraph selection and labeling, and human review until the F-1 score improves by less than ϵ .

The behavior of this algorithm is influenced by six parameters: N_0 , N , N_t , ϵ min, and max. N_0 and N are the number of paragraphs that are assigned to human reviewers in the first and subsequent steps, respectively. N_t is the number of examples in the test set. ϵ is a threshold that determines when to stop collecting data. The min and max determine the confidence range from which words are selected for human review. In the experimental studies described below, we used $N_0 = 278$, $N = 120$, $N_t = 500$, $\epsilon = 0$, min = 0.45, and max = 0.55.

The NER model used in the model-in-the-loop annotation workflow to score words might also be viewed as a parameter. In the work reported here, we use SpaCy exclusively for that purpose, as it integrates natively with the Prodigy annotation



tool and trains more rapidly. However, as we show below, the Keras-LSTM model is ultimately somewhat more accurate when trained on all of the labeled data generated, and thus is preferred when processing the entire CORPUS: see **Section 3.1.1** and **Section 3.2**.

This semi-automated method saves time and effort for human reviewers because they are only asked to verify labels that have already been identified by our model to be uncertain, and thus worth processing. Furthermore, as we show below, we find that we do not need to engage biomedical professionals to label drugs in text: untrained people, armed with contextual information (and online search engines), can spot drug names in text with accuracy comparable to that of experts.

We provide further details on the three phases of the algorithm in the following, with numbers in the list referring to line numbers in **Algorithm 1**.

A) Bootstrap

1. We start with the 2020–03–20 release version of the CORPUS, which contains 44,220 papers (Wang et al., 2020). We create \mathcal{C} , a random permutation of its paragraphs from which we will repeatedly fetch paragraphs via `next` (\mathcal{C}).
2. We bootstrap the labeling process by identifying as \mathcal{D} the 2,675 items in the DrugBank ontology with a Anatomical Therapeutic Chemical Classification System (ATC) code attached (eliminating many, but not all, drug-like molecule entities).
3. We create an initial set of silver paragraphs, \mathcal{P}_0 , by selecting N_0 paragraphs from \mathcal{C} that include a word from \mathcal{D} .

4. We engage human reviewers to remove false positives and label false negatives in \mathcal{P}_0 , yielding an initial set of gold paragraphs, \mathcal{B} .

B) Build test set

5. We expand the test set that we will use to evaluate the model created in the next phase, until we have N_t validated examples.
6. We train the NER model on 60% of the data collected to date and evaluate it on the remaining 40%, to create a new trained model, \mathcal{M} , with improved knowledge of the types of entities that we seek.
7. We use the probabilities over entities returned by the model to select, as our N new silver paragraphs, \mathcal{P} , paragraphs that contain at least one uncertain word (see above).
8. We engage human reviewers to convert these new silver paragraphs, \mathcal{P} , to gold, \mathcal{V} .
9. We add the new gold paragraphs, \mathcal{V} , to the bootstrap set \mathcal{B} .

- 11–12. Having assembled at least N_t validated examples, we select the first N_t as the test set, \mathcal{T} , and use any remaining examples to initialize the new gold set, \mathcal{G} .

C) Build labeled set

13. We assemble a training set \mathcal{G} , using the test set \mathcal{T} assembled in the previous phases for testing. This process continues until the F-1 score stops improving see **Section 2.2**.
- 14–17. Same as Steps 6–9, except that we train on \mathcal{G} and test on \mathcal{T} . Human reviewers are engaged to review new silver

Algorithm 1 | Model-in-the-loop Annotation Workflow

```

1  $\mathcal{C} := \text{randomized\_paragraphs}(\text{CORD})$ ; // Prepare CORD-19 paragraphs
  /* A) Bootstrap: Identify first silver paragraphs */
2  $\mathcal{D} := \{s : s \in \text{DrugBank} \ \& \ \text{ATC}(s)\}$ ; // Extract names from DrugBank
3  $\mathcal{P}_0 := \{p : p = \text{next}(\mathcal{C} \ \& \ p \cap \mathcal{D} \neq \emptyset) \ \& \ |\mathcal{P}_0| = N_0$ ; // Assemble first silver set
4  $\mathcal{B} = \text{verify}(\mathcal{P}_0)$ ; // Verify labels; initialize bootstrap set
  /* B) Build test set: Expand on bootstrap set from (A) */
5 while  $|\mathcal{B}| < N_t$  do
6    $\mathcal{M} := \text{NER}(\text{train} = 0.6\mathcal{B}, \text{test} = 0.4\mathcal{B})$ ; // Train: 60-40 train-test split
7    $\mathcal{P} := \{p : p = \text{next}(\mathcal{C} \ \& \ \exists w \in p : \mathcal{M}(w) \in [\text{min}, \text{max}])\} \ \& \ |\mathcal{P}| = N$ ; // Next silver
8    $\mathcal{V} := \text{verify}(\mathcal{P})$ ; // Engage crowd to verify labels
9    $\mathcal{B} := \mathcal{B} \cup \mathcal{V}$ ; // Add corrected paragraphs to bootstrap set
10 end
11  $\mathcal{T} = \mathcal{B}[N_t : ]$ ; // Initialize test set to first  $N_t$  examples
12  $\mathcal{G} = \mathcal{B}[N_t : ]$ ; // Initialize gold set to any remaining examples
  /* C) Build labeled set: Use test set from (B) to evolve model */
13 while  $F-1 \text{ score improvement} > \epsilon$  do
14    $\mathcal{M} := \text{NER}(\text{train} = \mathcal{G}, \text{test} = \mathcal{T})$ ; // Train on  $\mathcal{G}$ , test on  $\mathcal{T}$ 
15    $\mathcal{P} := \{p : p = \text{next}(\mathcal{C} \ \& \ \exists w \in p : \mathcal{M}(w) \in [\text{min}, \text{max}])\} \ \& \ |\mathcal{P}| = N$ ; // Next silver
16    $\mathcal{V} := \text{verify}(\mathcal{P})$ ; // Engage crowd to verify labels
17    $\mathcal{G} := \mathcal{G} \cup \mathcal{V}$ ; // Add corrected paragraphs to gold set
18 end

```

paragraphs and produce new gold paragraphs, which are then added to \mathcal{G} instead of \mathcal{T} .

2.2 Data-Performance Tradeoffs in Named-Entity Recognition Models

As noted in Section 2.1.2, our model-in-the-loop annotation workflow requires repeated retraining of a SpaCy model. Thus we conducted experiments to understand how SpaCy prediction performance is influenced by model size, quantity of training data, and amount of training performed.

As the training data produced by the model-in-the-loop evaluation workflow are to be used to train an NER model that we will apply to the entire CORD-19 dataset, we also evaluate the Keras-LSTM model from the perspectives of big data accuracy and training time.

2.2.1 Model Size

We first need to decide which SpaCy model to use for model-in-the-loop annotation. Model size is a primary factor that affects training time and prediction performance. In general, larger models tend to perform better, but require both more data and more time to train effectively. As our model-in-the-loop annotation strategy requires frequent model retraining, and furthermore will (initially at least) have little data, we hypothesize that a smaller model may be adequate for our purposes.

To explore this hypothesis, we study the performance achieved by the SpaCy medium and large models (Honnibal and Montani, 2020b) on our initial training set of 278 labeled paragraphs. We show in Figure 2 the performance achieved by the two models as a function of number of training epochs. Focusing on the harmonic mean of precision and recall, the F-1 score (a good measure a model's ability to recognize both true positives and true negatives), we see that the two models achieve similar prediction performance, with the largest difference in F-1 score being around 2%. As the large model takes over eight times longer to train per epoch, we select the medium model for model-in-the-loop data collection.

2.2.2 Word Embedding Models

The Keras LSTM model requires external word vectors since, unlike SpaCy, it does not include a word embedding model. To explore the affect of different word embedding models we trained both BERT (Devlin et al., 2018), a top-performing language model developed by Google, and FastText (Bojanowski et al., 2016), a model shown to have outperformed traditional Word2Vec models such as CBOW and Skipgram in our previous work (Hong et al., 2020b). While Google has released pre-trained BERT models, and researchers often build upon these models by "fine-tuning" them with additional training on small external datasets, it is not suitable to our problem as the vocabulary used in the CORD-19 is very different than the

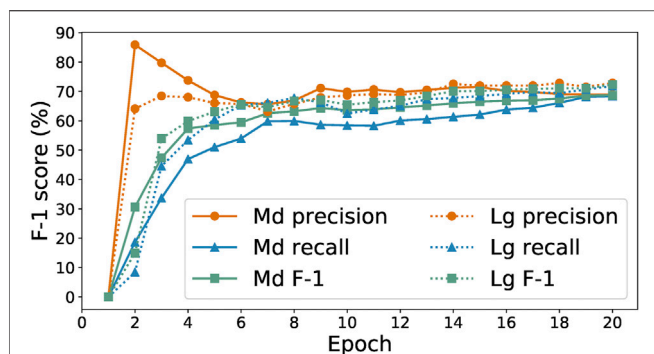


FIGURE 2 | Precision, recall, and F-1 scores of medium and large SpaCy models trained on 278 examples.

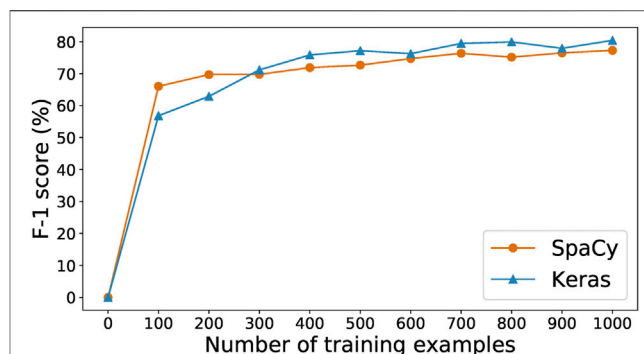


FIGURE 3 | Training curves for the SpaCy and Keras models for different number of examples collected.

datasets used to train these models.¹ Rather than use a pre-trained BERT model, we trained a BERT model on the CORD-19 corpus using a distributed neural network training framework from our previous work (Pauloski et al., 2020). As the CORD-19 corpus is approximately 20% of the size of the training data used by Google used to train BERT, we reduced the word embedding size proportionally from 768 to 128 to avoid over-fitting. For FastText word embeddings we set the size to the default 120.

We used word embeddings derived from both models to train the Keras LSTM model on the same training and testing data collected in **Section 2.1.2**. The model using FastText embeddings achieved a slightly higher F-1 score (80.5%) than the model trained with BERT embeddings (78.7%). This result is likely due to the limited training data and embedding size. In short, the humongous BERT model requires an equally humongous amount of data to achieve the best performance, and without such it will not necessarily outperform other much smaller and less computationally intensive word embedding models. In the remainder of the paper, we use the FastText word embedding model.

2.2.3 Amount of Training Data

As data labeling is expensive in both human time and model training time, it is valuable to explore the tradeoff between time spent collecting data and prediction performance. To this end, we manually labeled a set of 500 paragraphs selected at random from CORD-19 (Wang et al., 2020) as a test set. Then, we used that test set to evaluate the results of training the SpaCy and Keras-LSTM models of **Section 2.1.1** on increasing numbers of the paragraphs produced by our human-in-the-loop annotation process. **Figure 3** shows their F-1 score curves as we scale from 0 to 1,000 training samples. With only 100 training examples, SpaCy and Keras-LSTM achieve F-1 scores of 57 and 66%, respectively. SpaCy performs better than Keras-

LSTM with fewer training examples (i.e., less than 300), after which Keras-LSTM overtakes it and maintains a steady 2–3% advantage as the number of examples increases. This result motivates our choice of Keras-LSTM for the CORD-19 studies in **Section 3.2**.

We stopped collecting training data after 1,000 examples. We see in **Figure 3** that the performance of the SpaCy and Keras-LSTM models is essentially the same with 1,000 training examples as with 700 examples, with the F-1 score even declining when the number of available examples increases to 800 or 900. At 1,000 examples the F-1 score is greatest for both models. We conclude that the 1,000 training examples, along with the other 500 withheld as the test set, are best-suited to train our models. There are 4,244 and 1861 entities in the training and test set, respectively.

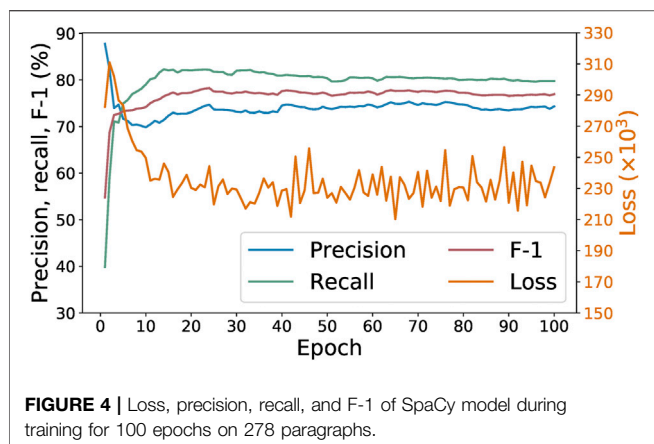
2.2.4 Training Epochs

Prediction performance is also influenced by the number of epochs spent in training. The cost of training is particularly important in a model-in-the-loop setup, as human reviewers cannot work while an model is offline for training.

Figure 4 shows the progression of the loss, precision, recall, and F-1 values of the SpaCy model during 100 epochs of training with the initial 278 examples. We can see that the best F-1 score is achieved within 10–20 epochs. Increasing the number of epochs does not result in any further improvement. Indeed, F-1 score does not tell us all about the model's performance. Sometimes training for more epochs could lead to lower loss values while other metrics (such as precision, recall, or F-1) no longer improve. That would still be desirable because it means the model is now more “confident,” in a sense, about its predictions. However, that is not the case here. As shown in **Figure 4**, after around 40 epochs the loss begins to oscillate instead of continuing downwards, suggesting that in this case training for 100 epochs does not result in a better model than only training for 20 epochs.

Figure 5 shows the progression of accuracy and loss value for the Keras-LSTM model with the initial 278 examples. In **Figure 5A**, we see that validation accuracy improves as training accuracy increases during the first 50 epochs. After around epoch 50, the training and validation accuracy curves

¹A biomedical-focused BERT model PubMdBERT (Gu et al., 2020) was released after the work described in this paper was done. We have fine-tuned the PubMedBERT-base-uncased-abstract-fulltext model on our dataset with the following parameters: batch size: 16, warmup steps: 500, weight decay: 0.01, training epochs: 3. It achieved a F-1 score of 0.885.



diverge: the training accuracy continues to increase but the validation accuracy plateaus. This trend is suggestive of overfitting, which is corroborated by **Figure 5B**. After about 50 epochs, the validation loss curve turns upwards. Hence we choose to limit the training epochs to 64. After each epoch, if a lower validation loss is achieved, the current model state is saved. After 64 epochs, we test the model with the lowest validation loss on the withheld test set.

3 RESULTS

We present the results of experiments in which we first evaluate the performance of our models from various perspectives and then apply the models to the **CORD-19** dataset.

3.1 Evaluating Model and Human Performance

We conducted experiments to compare the performance of the SpaCy and Keras-LSTM NER models; compare the performance of the models against humans; determine how training data influences model performance; and analyze human and model errors.

3.1.1 Performance of SpaCy and Keras Named-Entity Recognition Models

We used the collected data of **Section 2.1.2** to train both the SpaCy and Keras-LSTM NER models of **Section 2.1.1** to recognize and extract drug-like molecules in text. The SpaCy model used is the medium English language model `en_core_web_md` (Honnibal and Montani, 2020b). For the Keras model, the input embedding size is 120, the LSTM and Fully-connected hidden layers have a size of 32, and the dropout rate is 0.1. The model is trained for 64 epochs with a batch size of 64. We find that the trained SpaCy model achieved a best F-1 score of 77.3%, while the trained Keras-LSTM model achieved a best F-1 score of 80.5%, somewhat outperforming SpaCy.

As shown in **Figure 3**, the SpaCy model performs better than the Keras-LSTM model when trained with small amounts of training data—perhaps because of the different mechanisms employed by the two methods to generate numerical representations for words. SpaCy's built-in language model, pre-trained on a general corpus of blog posts, news, comments, etc., gives it some knowledge about commonly used words in English, which are likely also to appear in a scientific corpus. On the other hand, the Keras-LSTM model uses custom word embeddings trained solely on an input corpus, which provides it with better understanding of multi-sense words, especially those that have quite different meanings in a scientific corpus. However, without enough raw data to draw contextual information from, custom word embeddings can not accurately reflect the meaning of words.

3.1.2 Comparison Against Human Performance

Recognizing drug-like molecules is a difficult task even for humans, especially non-medical professionals (such as our non-expert annotators). To assess the accuracy of the annotators, we asked three people to examine 96 paragraphs, with their associated labels, selected at random from the labeled examples. Two of these reviewers had been involved in creating the labeled dataset; the third had not. For each paragraph, each reviewer decided independently whether each drug molecule entity was labeled correctly (a true positive), was labeled as a

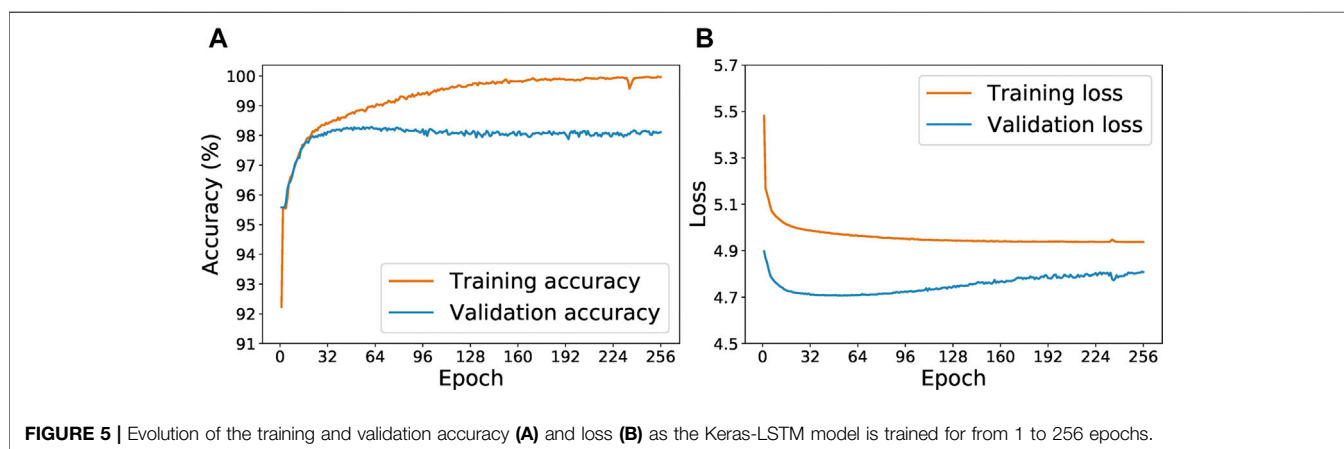


TABLE 1 | K-fold ($K = 5$) validation of the SpaCy model on 96 paragraphs with original vs. corrected labels. The first five rows are the results of each fold; the last row is the average F-1 score of the five folds.

Original labels			Corrected labels		
Precision	Recall	F-1	Precision	Recall	F-1
100.0	25.6	40.7	100.0	20.9	34.6
93.8	50.6	65.7	88.9	53.9	67.1
93.0	63.5	75.5	92.0	73.0	81.4
76.5	45.2	57.1	68.6	61.4	64.8
98.9	92.3	95.5	99.2	92.1	95.5
Average	—	66.9	Average	—	68.7

drug when it was not (a false positive), or was not labeled (a false negative). If all three reviewers agreed in their opinions on a paragraph (the case for 88 of the 96 paragraphs), we accepted their opinions; if they disagreed (the case for eight paragraphs), we engaged an expert.

This process revealed a total of 257 drug molecule entities in the 96 paragraphs, of which the annotators labeled 201 correctly (true positives), labeled 49 incorrectly (false positives), and missed 34 (false negatives). The numbers of true positives and false negatives do not sum up to the total number of drug molecules because in some cases an annotator labeled not to a drug entity but the entity plus extra preceding or succeeding word or punctuation mark (e.g., “sofosbuvir,” instead of “sofosbuvir”) and we count such occurrences as false positives rather than false negatives. In this evaluation, the non-expert annotators achieved an F-1 score of 82.9%, which is comparable to the 80.5% achieved by our automated models, as shown in **Figure 3**. In other words, our models have performance on par with that of non-expert humans.

3.1.3 Effects of Training Data Quality on Model Performance

We described in the previous section how review of 96 paragraphs labeled by the non-expert annotators revealed an error rate of about 20%. This raises the question of whether model performance could be improved with better training data. To examine this question, we compare the performance of our models when trained on original vs. corrected data. As we only have 96 corrected paragraphs, we restrict our training sets to those 96 paragraphs in each case.

We sorted the 96 paragraphs in both datasets so that they are considered in the same order. Then, we split each dataset into five subsets for K-fold cross validation ($K = 5$), with the first four subsets having 19 paragraphs each and the last subset having 20. Since K is set to five, the SpaCy and Keras models are trained five times. In the i th round, each model is trained on four subsets (excluding the i th) of each dataset. The i th subset of the corrected dataset is used as the test set. The i th subset of the original dataset is not used in the i th round.

We present the K-fold cross validation results in **Tables 1, 2**. The models performed reasonably well when trained on the

TABLE 2 | K-Fold ($K = 5$) validation of the Keras-LSTM model on 96 paragraphs with original vs. corrected labels. The first five rows are the results of each fold; the last row is the average F-1 score of the five folds.

Original Labels			Corrected Labels		
Precision	Recall	F-1	Precision	Recall	F-1
88.5	53.5	66.7	80.6	69.1	74.4
70.6	57.8	63.6	90.2	59.1	71.4
70.0	67.7	68.9	82.9	47.5	60.4
80.8	55.3	65.6	80.8	51.2	62.7
78.1	56.8	65.8	75.0	67.9	71.3
Average	—	66.1	Average	—	68.0

original dataset, with an average F-1 score only 2% less than that achieved with the corrected labels. Given that the expert input required for validation is hard to come by, we believe that using non-expert reviewers is an acceptable tradeoff and probably the only practical way to gather large amounts of training data.

3.2 Applying the Trained Models

After training the models with the labeled examples, we applied the trained models to the entire CORON-19 corpus (2020–10-04 version with 198,875 articles) to identify potential drug-like molecules. Processing a single article takes only a few seconds; we adapted our models to use data parallelism to enable rapid processing of these many articles.

We ran the SpaCy model on two Intel Skylake 6,148 processors with a total of 40 CPU cores; this run took around 80 core-hours and extracted 38,472 entities. We ran the Keras model on four NVidia Tesla V100 GPUs; this run took around 40 GPU-hours and extracted 121,680 entities. We recorded for each entity the number of the times that it has been recognized by each model, and used those numbers as a voting mechanism to further determine which entities are the most likely to be actual drugs. In our experiments, “balanced” entities (i.e., those whose numbers of detection by the two models are within a factor of 10 of each other) are most likely to appear in the DrugBank list. As shown in **Figure 6**, we sorted all extracted entities in descending order by their total number of detections by both models and compared the top 100 entities to DrugBank. We found that only 77% were exact matches to drug names or aliases, or 86% if we included partial matches (i.e., the extracted entity is a word within a multi-word drug name or alias in DrugBank). In comparison, among the top 100 “balanced” entities, 88% were exact matches to DrugBank, or 91% with partial matches.

Although DrugBank provides a reference metric to evaluate the results, it is not an exhaustive list of known drugs. For instance, remdesivir, a drug that has been proposed as a potential cure for COVID-19, is not in DrugBank. We manually checked via Google searches the top 50 “balanced” and top 50 “imbalanced” entities not matched to DrugBank, and found that 70% in the “balanced” list are actual drugs, but only 26% in the “imbalanced” list. Looking at the false positives in these top 50 lists, the “balanced” false positives are often

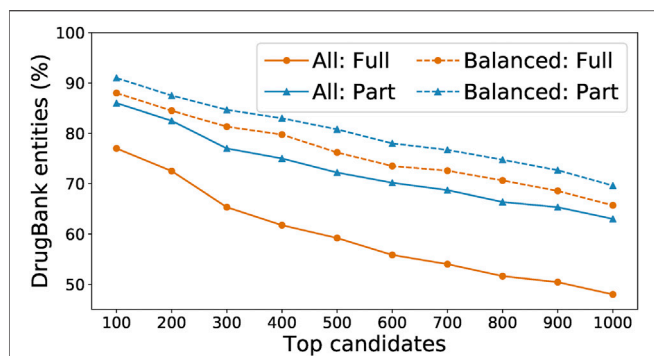


FIGURE 6 | Percentage of detected entities that are also found in DrugBank, when running either on all words found by our model or on just the balanced subset, and with “found” defined as either a full or partial match.

understandable. For example, in the sentence “ELISA plate was coated with ... and then treated for 1 h at 37.8°C with dithiothreitol ...”, the model mistook the redox reagent *dithiothreitol* for a drug entity, probably due to its context “treated with.” On the other hand, we found no such plausible explanations for the false positives in the “imbalanced” list, where most false positives are chemical elements (e.g., silver, sodium), amino acids (e.g., cysteine, glutamine), or proteins (e.g., lactoferrin, cystatin).

Finally, we compared our extraction results to the drugs being used in clinical trials, as listed on the United States National Library of Medicine website (National Institutes of Health (2020)). We queried the website with “covid” as the keyword and manually screened the returned drugs in the “Interventions” column to remove stopwords (e.g., tablet, injection, capsule) and dosage information (e.g., 2.5 mg, 2.5%) and only keep the drug names. Then we compared the top 50 most frequently appeared drugs to the automatically extracted drugs from literature. The “balanced” entities extracted by both models matched to 64% of the top 50 drugs in clinical trial, whereas the “imbalanced” entities only matched to 6% in the same list.

The results discussed here are available in the repository described in **Section 5**.

3.3 Validating the Utility of Identified Molecules

We are interested to evaluate the relevance to COVID research of the molecules that we extracted from CORD19. To that end, we compared our extracted molecule list against ZINC and Drugbank sets that a group of scientists at Argonne National Laboratory had used for computational screening. We found that our list contained an additional 3,591 molecules not found in their screening sets (filtered by their canonical SMILE strings). Applying their methods to screen those 3,591 molecules for docking against a main coronavirus protease, 3CLPro, revealed that 18 had docking scores in the top 0.1% of the 6.6 M ZINC molecules that they had screened previously (Saadi et al., 2020)—significantly more than the four that we would expect by chance.

TABLE 3 | The 10 most frequent tokens, excluding stopwords and punctuation marks, within various window sizes around entities incorrectly labeled by human reviewers.

#	Window size = 1		Window size = 3		Window size = 5	
	Token	Count	Token	Count	Token	Count
1	300	4	Mg	19	mg	23
2	oral	3	once/day	7	daily	8
3	dose	3	treatment	6	treatment	8
4	intravenous	2	300	5	once/day	7
5	500	2	treated	4	300	7
6	intravenously	1	Oral	4	oral	6
7	include	1	Once	4	recipients	5
8	Both	1	Dose	4	treated	4
9	resistance	1	Cidofovir	3	twice	4
10	treatment	1	resistance	3	include	4

As reported by Babuji et al. (2020a), those researchers have leveraged the outputs of our models in a computational screening pipeline that leverages HPC resources at scale, coupled with multiple artificial intelligence and simulation-based approaches (including leads from NLP), to identify high-quality therapeutic compounds to screen experimentally. A further manuscript, detailing the end-to-end process from data collection to simulation, incorporation of these results, cellular assays, and identification of high performing therapeutic compounds, is in preparation.

4 DISCUSSION: ANALYSIS OF HUMAN AND MODEL ERRORS

Finally, we explore the contexts in which human reviewers and models make mistakes. Specifically, we study the tokens that appear most frequently near to incorrectly labeled entities. To investigate the effects of immediate and long-distance context, we control, as *window size*, the maximum distance between a token and a entity for that token to be considered as “context” for that entity.

One difficulty with this analysis is that the most frequent tokens identified in this way were mostly stop words or punctuation marks. For instance, when the window size is set to three, the 10 most frequent tokens around mislabeled words are, in descending order, “comma (,)”, “and,” “mg,” “period (.),” “right parenthesis (),” “with,” “of,” “left parenthesis ((),” “is,” and “or.” Only “mg” is neither a stop word nor punctuation mark.

Those tokens provide little insight as to why human reviewers might have made mistakes, and furthermore are unlikely to have influenced reviewer decisions. Thus we exclude stopwords and punctuation marks when providing, in **Table 3**, lists of the 10 most frequent tokens within varying window sizes of words that were *incorrectly* identified as molecules by human reviewers.

TABLE 4 | The 10 most frequent tokens, excluding stop words and punctuation marks, within various window sizes around entities correctly labeled by human reviewers.

#	Window size = 1		Window size = 3		Window size = 5	
	Token	Count	Token	Count	Token	Count
1	resistance	176	Tetracycline	230	Tetracycline	230
2	treatment	9	resistance	177	resistance	178
3	mM	4	Trimethoprim	118	Trimethoprim	118
4	oral	3	treatment	11	treatment	14
5	after	3	20 ~	7	20 ~	8
6	analogue	3	Figure	5	placebo	7
7	responses	3	concentration	5	effects	6
8	antibiotics	2	compared	4	Figure	6
9	exposure	2	100	4	KLK5	6
10	pharmacokinetics	2	mM	4	Matriptase	6

TABLE 5 | The 20 most frequent tokens, including stop words and punctuation marks, within various window sizes around entities incorrectly labeled by human reviewers. Words that are neither stop words nor punctuation words are in boldface.

#	Window size = 1		Window size = 3		Window size = 5	
	Token	Count	Token	Count	Token	Count
1	—	27	—	49	—	74
2	—	14	And	21	—	28
3	And	14	Mg	19	And	28
4	with	8	—	18	Mg	23
5	(7)	13)	21
6	Is	7	With	10	Of	18
7	Of	6	Of	10	(17
8	Was	6	(9	with	13
9	Or	4	Is	9	to	12
10	300	4	Or	7	the	11
11	oral	3	once/day	7	a	11
12	Has	3	A	7	in	10
13	To	3	The	6	is	9
14	[3	Was	6	or	9
15	dose	3	To	6	daily	8
16	intravenous	2	treatment	6	was	8
17	In	2	In	5	treatment	8
18	may	2	300	5]	8
19	500	2	Be	4	once/day	8
20	A	2	Treated	4	were	7

TABLE 6 | The 20 most frequent tokens, including stop words and punctuation marks, within various window sizes around entities incorrectly labeled by the Keras model. Words that are neither stop words nor punctuation words are in boldface.

#	Window size = 1		Window size = 3		Window size = 5	
	Token	Count	Token	Count	Token	Count
1	—	166	—	347	—	468
2	(86	and	126	and	176
3	And	81	(117	(162
4	Of	58)	89)	143
5)	30	of	85	of	130
6	To	28	the	73	the	121
7	Or	24	to	60	to	85
8		22		48	in	75
9	mM	18	with	44	with	72
10	With	17	a	41		68
11	In	15	in	37	a	62
12	For	15	or	35	was	52
13	Is	14	was	33	or	47
14	As	14	mM	32	is	43
15	The	13	is	31	for	42
16	Was	12	as	29	that	37
17	That	11	that	25	mM	36
18	[11	by	22	by	35
19	treatment	9	for	22	as	32
20	A	9	[22	were	31

We see that there are indeed several deceptive contextual words. With a window size of one, the 10 most frequent tokens include “oral,” “dose,” and “intravenous.” It is understandable that an untrained reviewer might label as drugs words that immediately precede or follow such context words. Similar patterns can be seen for window sizes of three and five. Without background knowledge to draw from, non-experts are more likely to rely on their experience gained from labeling previous paragraphs. One may hypothesize that after the reviewers have seen a few dozen to a few hundred paragraphs, those deceptive contextual words must have left a deep impression, so that when those words re-appear they are likely to label the strange unknown word close to them as a drug.

To investigate this hypothesis, we also explored the most frequent words around drug entities that are *correctly* labeled by human reviewers: see **Table 4**. Interestingly, we found overlaps between the lists in **Tables 3, 4**: in all, three, four, and two overlaps for window sizes of one, three, and five, respectively, when treating all numerical values as identical. This finding supports our hypothesis that those frequent words around real drug entities may confuse human reviewers when they appear around non-drug entities.

We repeat this comparison of context words around human and model errors while considering stopwords and punctuation marks. **Tables 5, 6** show the 20 most frequent tokens in each case. We see that 20–25% of the tokens in **Table 5**, but only 5–10% of those in **Table 6**, are *not* stop words or punctuation marks. As the model only learns its word embeddings from the input text, if a token often co-occurs with drug entities in the training corpus the model will treat it as an indication of drug entities near its presence, regardless of whether or not it is a stopword. This apparently leads the model to make incorrect inferences. Humans, on the other hand, are unlikely to think that stopword such as “the” is indicative of drug entities, no matter how frequently they appear together.

5 DATA AVAILABILITY AND FORMATS

We have made our annotated training data, trained models, and the results of applying the models to the CORD-19 corpus publicly available online. (Babuji et al., 2020b).

In order to facilitate training of various models, we published the training data in two formats—an unsegmented version in line-delimited JSON (JSONL) format, and a segmented version in Comma Separated Value (CSV) format. The JSONL format contains the most comprehensive information that we have collected on the paragraphs in the dataset. We choose JSONL format rather than a JSON list because it allows for the retrieval of objects without having to parse the entire file. A JSON object in the JSONL file has the following structure:

- text: The original paragraph stored as a string without any modification.
- tokens: The list of tokens from text after tokenization.
- text: The text of the token as a string.
- start: The index of the first character of the token in text.
- end: The index of the first character after the token in text.
- id: Zero-based numbering of the token.

- spans: The list of spans (sequences of tokens) that are labeled as named entities (drugs)
- start: The index of the first character of the span in text.
- end: The index of the first character after the span in text.
- token_start: The index of the first token of the span in text.
- token_end: The index of the last token of the span in text.
- label: The label of the span (“drug”)

Another commonly adopted labeling scheme for NER datasets is the “IOB” labeling scheme, in which the original text is first tokenized and each token is assigned a label “I,” “O,” or “B.” The label “B (eginning)” means the corresponding token is the first in a named entity. A label “I (nside)” is given to every token in a named entity except for the first token. All other tokens gets the label “O (outside)” which means they are not part of any named entity. The aforementioned JSONL data are converted according to the IOB scheme and stored in Comma Separated Value (CSV) files with one training example per line. Each line consists of two columns: a first of tokens that made up of the original texts, and a second of the corresponding IOB labels for those tokens. In addition to a different labeling scheme, the samples in the CSV files are segmented, meaning that each sentence is treated as a training sample instead of an entire paragraph. This structure aligns with that used in standard NER training sets such as CoNLL03 (Sang and De Meulder, 2003).

The trained SpaCy and Keras models and the results of applying the models to the 198,875 articles in the CORD-19 corpus are also available in this GitHub repo. Additionally, the pre-trained SpaCy model is provided as a cloud service via DLHub (Hong et al., 2020a; Li et al., 2021). (The Keras model could not be hosted there due to compatibility issues with DLHub.) This cloud service allows researchers to apply the model to any texts they provide with as few as four lines of code.

6 CONCLUSION AND FUTURE DIRECTIONS

We have presented a human-machine hybrid pipeline for collecting training data for named entity recognition (NER) models. We applied this pipeline to create a NER model for identifying drug-like molecules in COVID-19-related research papers. Our pipeline facilitated efficient use of valuable human resources by presenting human labellers only with samples that were most likely to confuse the NER model. We explored various trade-offs, including model size, number of training samples, and training epochs, to find the right balance between model performance and time-to-result. In total, human reviewers working with our pipeline validated labels for 278 bootstrap samples, 1,000 training samples, and 500 test samples. As this work was performed in conjunction with other tasks, we cannot accurately quantify the total effort taken to collect and annotate the above training and test samples, but it was likely around 100 person-hours.

NER models trained with these training data achieved a best F-1 score of 80.5% when evaluated on our collected test set. Our models

correctly identified 64% of the top 50 drugs that are in clinical trials for COVID-19, and when applied to 198,875 articles in the COVID-19 collection, identified 10,912 molecules with potential therapeutic effects against the SARS-CoV-2 coronavirus. The code, model, and extraction results are publicly available. Our work provided an additional 3591 SMILES strings to scientists at Argonne National Laboratory to be used in computational screening pipelines, of which 18 ranked in the top 0.1% of the molecules screened. Babuji et al. (2020a) have leveraged the outputs of our models in a computational screening pipeline that leverages HPC resources at scale to identify high-quality therapeutic compounds to screen experimentally. A further manuscript detailing the end-to-end process of identifying high performing therapeutic compounds is in preparation.

DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/globus-labs/covid-nlp/tree/master/drug-ner>.

AUTHOR CONTRIBUTIONS

ZH and JGP conducted the computational experiments. BB and LW contributed to discussions of ideas. ZH wrote the first draft.

REFERENCES

- Allen Institute For AI (2020). [Dataset]: COVID-19 Open Research Dataset challenge. <https://www.kaggle.com/allen-institute-for-ai/CORD-19-research-challenge>.
- Babuji, Y., Blaiszik, B., Brettin, T., Chard, K., Chard, R., Clyde, A., et al. (2020a). *Targeting SARS-CoV-2 with AI-And HPC-Enabled lead Generation: A First Data Release*. arXiv. preprint arXiv:2006.02431.
- Babuji, Y., Blaiszik, B., Chard, K., Chard, R., Foster, I., Gordon, I., et al. (2020b). [Dataset]: Lit - A Collection of Literature Extracted Small Molecules to Speed Identification of COVID-19 Therapeutics. Materials Data Facility. <https://github.com/globus-labs/covid-nlp/tree/master/drug-ner>.
- Bada, M., Eckert, M., Evans, D., Garcia, K., Shipley, K., Sitnikov, D., et al. (2012). Concept Annotation in the CRAFT Corpus. *BMC bioinformatics*. 13, 161. doi:10.1186/1471-2105-13-161
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2016). Enriching Word Vectors with Subword Information. *Trans. Assoc. Comput. Linguistics*. 5, 135–146.
- Bullard-Feibelman, K. M., Govero, J., Zhu, Z., Salazar, V., Veselinovic, M., Diamond, M. S., et al. (2017). The FDA-Approved Drug Sofosbuvir Inhibits Zika Virus Infection. *Antiviral Res.* 137, 134–140. doi:10.1016/j.antiviral.2016.11.023
- Center for Drug Evaluation and Research (2020). [Dataset]: About Drugs@FDA. <https://www.fda.gov/drugs/drug-approvals-and-databases/about-drugs@fda>.
- Chiu, J. P. C., and Nichols, E. (2016). Named Entity Recognition with Bidirectional LSTM-CNNs. *TACL* 4, 357–370. doi:10.1162/tACL_a_00104
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. arXiv. preprint arXiv:1810.04805.
- Dong, E., Du, H., and Gardner, L. (2020). An Interactive Web-Based Dashboard to Track COVID-19 in Real Time. *Lancet Infect. Dis.* 20, 533–534. doi:10.1016/s1473-3099(20)30120-1
- KC and IF revised the manuscript. All authors contributed to the article and approved the submitted version.

FUNDING

This research was supported by the DOE Office of Science through the National Virtual Biotechnology Laboratory, a consortium of DOE national laboratories focused on response to COVID-19, with funding provided by the Coronavirus CARES Act. This research used resources of the Argonne Leadership Computing Facility, a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. This work was also supported by financial assistance award 70NANB19H005 from United States Department of Commerce, National Institute of Standards and Technology as part of the Center for Hierarchical Materials Design (CHiMaD).

ACKNOWLEDGMENTS

The data generated have been prepared as part of the nCov-Group Collaboration, a group of over 200 researchers working to use computational techniques to address various challenges associated with COVID-19. We are grateful to our human reviewers: India S. Gordon, Linda Novak, Kasia Salim, Susan Sarvey, Julie Smagacz, and Monica Orozco White.

- Furrer, L., Jancso, A., Colic, N., and Rinaldi, F. (2019). OGER++: Hybrid Multi-type Entity Recognition. *J. Cheminform.* 11, 7–10. doi:10.1186/s13321-018-0326-3
- Gu, Y., Tinn, R., Cheng, H., Lucas, M., Usuyama, N., Liu, X., et al. (2020). *Domain-specific Language Model Pretraining for Biomedical Natural Language Processing*. arXiv. preprint arXiv:2007.15779.
- Hong, Z., Pauloski, J., Ward, L., Chard, K., Blaiszik, B., and Foster, I. (2020a). [Dataset]: A Model to Extract Drug-like Molecules from Natural Language Text. Data and Learning Hub for Science. <https://10.26311/EV06-C385>.
- Hong, Z., Tchoua, R., Chard, K., and Foster, I. (2020b). “SciNER: Extracting Named Entities from Scientific Literature,” in *International Conference on Computational Science* (Amsterdam: Springer), 308–321. doi:10.1007/978-3-030-50417-5_23
- Honnibal, M., and Montani, I. (2020a). [Dataset]: SpaCy Pipeline Design. <https://spacy.io/models#design>.
- Honnibal, M., and Montani, I. (2020b). [Dataset]: SpaCy Pretrained English Models. <https://spacy.io/models/en>.
- Honnibal, M., Montani, I., Van Landeghem, S., and Boyd, A. (2020). [Dataset]: spaCy: Industrial-Strength Natural Language Processing in Python. <https://10.5281/zenodo.1212303>
- Li, Z., Chard, R., Ward, L., Chard, K., Skluzacek, T. J., Babuji, Y., et al. (2021). DLHub: Simplifying Publication, Discovery, and Use of Machine Learning Models in Science. *J. Parallel Distributed Comput.* 147, 64–76. doi:10.1016/j.jpdc.2020.08.006
- Nadeau, D., and Sekine, S. (2007). A Survey of Named Entity Recognition and Classification. *Lingvist. Investig.* 30, 3–26. doi:10.1075/li.30.1.03nad
- National Institutes of Health (2020). [Dataset]: U.S. COVID-19 Clinical Trials. <https://clinicaltrials.gov/ct2/results?cond=covid&cntry=US>.
- Oestereich, L., Rieger, T., Neumann, M., Bernreuther, C., Lehmann, M., Krasemann, S., et al. (2014). Evaluation of Antiviral Efficacy of Ribavirin, Arbidol, and T-705 (Favipiravir) in a Mouse Model for Crimean-Congo Hemorrhagic Fever. *Plos Negl. Trop. Dis.* 8, e2804. doi:10.1371/journal.pntd.0002804
- Pauloski, J., Zhang, Z., Huang, L., Xu, W., and Foster, I. (2020). “Convolutional Neural Network Training with Distributed K-FAC,” in 2020 SC20: International Conference

- for High Performance Computing, Networking, Storage and Analysis (SC) (Atlanta, GA: IEEE Computer Society, 1331–1342. doi:10.1109/sc41405.2020.00098
- Rindflesch, T. C., Tanabe, L., Weinstein, J. N., and Hunter, L. (1999). EDGAR: Extraction of Drugs, Genes and Relations from the Biomedical Literature. *Biocomputing (World Scientific)* 2000, 517–528.
- Saadi, A. A., Alfe, D., Babuji, Y., Bhati, A., Blaiszik, B., Brettin, T., et al. (2020). *IMPECCABLE: Integrated Modeling Pipeline for COVID Cure by Assessing Better Leads*. arXiv. preprint arXiv:2010.06574.
- Sang, E. F., and De Meulder, F. (2003). “Introduction to the CoNLL-2003 Shared Task: Language-independent Named Entity Recognition,” in 7th Conference on Natural Language Learning Edmonton, 142–147.
- Tchoua, R., Ajith, A., Hong, Z., Ward, L., Chard, K., Audus, D., et al. (2019a). “Active Learning Yields Better Training Data for Scientific Named Entity Recognition,” in 15th International Conference on eScience, San Diego, CA (Faro, Algarve: IEEE), 126–135. doi:10.1109/escience.2019.00021
- Tchoua, R. B., Ajith, A., Hong, Z., Ward, L. T., Chard, K., Belikov, A., et al. (2019b). “Creating Training Data for Scientific Named Entity Recognition with Minimal Human Effort,” in *International Conference on Computational Science*. Springer, 398–411. doi:10.1007/978-3-030-22734-0_29
- Wang, L. L., Lo, K., Chandrasekhar, Y., Reas, R., Yang, J., Eide, D., et al. (2020). *CORD-19: The COVID-19 Open Research Dataset*. ArXiv. Preprint 2004.10706.
- Weischedel, R., Palmer, M., Marcus, M., Hovy, E., Pradhan, S., Ramshaw, L., et al. (2013). *OntoNotes Release 5.0*. Philadelphia, PA: Linguistic Data Consortium, 23.
- Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., et al. (2018). DrugBank 5.0: A Major Update to the DrugBank Database for 2018. *Nucleic Acids Res.* 46, D1074–D1082. doi:10.1093/nar/gkx1037

Author Disclaimer: Unless otherwise indicated, this information has been authored by an employee or employees of the UChicago Argonne, LLC, operator of the Argonne National laboratory with the United States Department of Energy. The United States Government has rights to use,

reproduce, and distribute this information. The public may copy and use this information without charge, provided that this Notice and any statement of authorship are reproduced on all copies.

While every effort has been made to produce valid data, by using this data, User acknowledges that neither the Government nor UChicago Argonne LLC makes any warranty, express or implied, of either the accuracy or completeness of this information or assumes any liability or responsibility for the use of this information. Additionally, this information is provided solely for research purposes and is not provided for purposes of offering medical advice. Accordingly, the United States Government and UChicago Argonne LLC are not to be liable to any user for any loss or damage, whether in contract, tort (including negligence), breach of statutory duty, or otherwise, even if foreseeable, arising under or in connection with use of or reliance on the content displayed on this site.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations or those of the publisher, the editors, and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Hong, Pauloski, Ward, Chard, Blaiszik and Foster. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.