

Ontology-driven weak supervision for clinical entity classification in electronic health records

Jason A. Fries ¹✉, Ethan Steinberg^{1,2}, Saelig Khattar², Scott L. Fleming ¹, Jose Posada ¹, Alison Callahan ¹ & Nigam H. Shah ¹

In the electronic health record, using clinical notes to identify entities such as disorders and their temporality (e.g. the order of an event relative to a time index) can inform many important analyses. However, creating training data for clinical entity tasks is time consuming and sharing labeled data is challenging due to privacy concerns. The information needs of the COVID-19 pandemic highlight the need for agile methods of training machine learning models for clinical notes. We present Trove, a framework for weakly supervised entity classification using medical ontologies and expert-generated rules. Our approach, unlike hand-labeled notes, is easy to share and modify, while offering performance comparable to learning from manually labeled training data. In this work, we validate our framework on six benchmark tasks and demonstrate Trove's ability to analyze the records of patients visiting the emergency department at Stanford Health Care for COVID-19 presenting symptoms and risk factors.

¹Center for Biomedical Informatics Research, Stanford University, Stanford, CA, USA. ²Department of Computer Science, Stanford University, Stanford, CA, USA. ✉email: jason-fries@stanford.edu

Analyzing text to identify concepts such as disease names and their associated attributes like negation are foundational tasks in medical natural language processing (NLP). Traditionally, training classifiers for named entity recognition (NER) and cue-based entity classification have relied on hand-labeled training data. However, annotating medical corpora requires considerable domain expertise and money, creating barriers to using machine learning in critical applications^{1,2}. Moreover, hand-labeled datasets are static artifacts that are expensive to change. The recent COVID-19 pandemic highlights the need for machine-learning tools that enable faster, more flexible analysis of clinical and scientific documents in response to rapidly unfolding events³.

To address the scarcity of hand-labeled training data, machine-learning practitioners increasingly turn to lower cost, less accurate label sources to rapidly build classifiers. Instead of requiring hand-labeled training data, weakly supervised learning relies on task-specific rules and other imperfect labeling strategies to programmatically generate training data. This approach combines the benefits of rule-based systems, which are easily shared, inspected, and modified, with machine learning, which typically improves performance and generalization properties. Weakly supervised methods have demonstrated success across a range of NLP and other settings^{4–8}.

Knowledge bases and ontologies provide a compelling foundation for building weakly supervised entity classifiers. Ontologies codify a vast amount of medical knowledge via taxonomies and example instances for millions of medical concepts. However, repurposing ontologies for weak supervision creates challenges when combining label information from multiple sources without access to ground-truth labels. The hundreds of terminologies found in the Unified Medical Language System (UMLS) Metathesaurus⁹ and other sources¹⁰ typify the highly redundant, conflicting, and imperfect entity definitions found across medical ontologies. Naively combining such conflicting label assignments can cause substantial performance drops in weakly supervised classification¹¹; therefore, a key challenge is correcting for labeling errors made by individual ontologies when combining label information.

Rule-based systems for NER and cue detection^{12,13} are common in clinical text processing, where labeled corpora are difficult to share due to privacy concerns. Generating imperfect training labels from indirect sources (e.g., patient notes) is often used in analyzing medical images^{14–16}. Recent work has explored learning the accuracies of sources to correct for label noise when using rule-based systems to generate training data for text

classification^{4,17}. Weakly supervised clinical applications have explored document and relation classification using task-specific rules^{18,19} or leveraging dependency parsing and compositional grammars to automate relation classification for standardizing clinical concepts²⁰. However, these largely focus on relation and document classification via task-specific labeling rules or sourcing supervision from a single ontology and do not explore NER or automating labeling via multiple ontologies.

Prior research on weakly supervised NER has required complex preprocessing to identify possible entity spans²¹, generated labels from a single source rather than combining multiple sources²², or relied on ad hoc rule engineering²³. High-impact application areas, such as clinical NER using weak supervision, are largely unstudied. Recent weak supervision frameworks such as Snorkel¹¹ are domain- and task-agnostic, introducing barriers to quickly developing and deploying labeling heuristics in complex domains such as medicine. Key questions remain about the extent to which we can automate weak supervision using existing medical ontologies and how much additional task-specific rule engineering is required for state-of-the-art performance. It is also unclear whether, and by how much, pretrained language models such as BioBERT²⁴ improve the ability to generalize from weakly labeled data and reduce the need for task-specific labeling rules.

We present a Trove, a framework for training weakly supervised medical entity classifiers using off-the-shelf ontologies as a source of reusable, easily automated labeling heuristics. Doing so transforms the work of using weak supervision from that of coding task-specific labeling rules to defining a target entity type and selecting ontologies with sufficient coverage for a target dataset, which is a common interface for popular biomedical annotation tools such as NCBO BioPortal and MetaMap^{10,25}. We examine whether ontology-based weak supervision, coupled with recent pretrained language models such as BioBERT, reduces the engineering cost of creating entity classifiers while matching the performance of prior, more expensive, weakly supervised approaches. We further investigate how ontology-based labeling functions can be extended when we need to incorporate additional, task-specific rules. The overall pipeline is shown in Fig. 1.

In this work, we demonstrate the utility of Trove through six benchmark tasks for clinical and scientific text, reporting state-of-the-art weakly supervised performance (i.e., using no hand-labeled training data) on NER datasets for chemical/disease and drug tagging. We further present weakly supervised baselines for two tasks in clinical text: disorder tagging and event temporality classification. Using ablation analyses, we characterize the performance trade-offs of training models with labels generated from

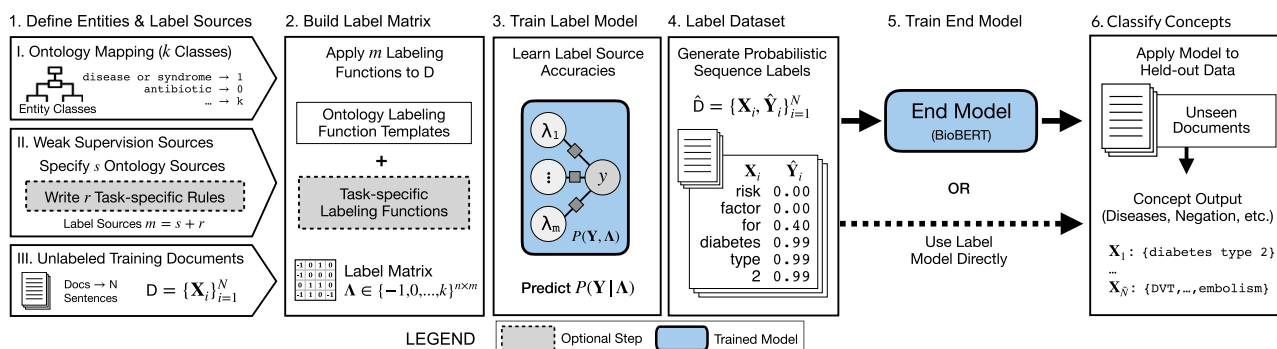


Fig. 1 Trove pipeline for ontology-driven weak supervision for medical entity classification. Users specify (I) a mapping of an ontology's class categories to entity classes, (II) a set of label sources (e.g., ontologies, task-specific rules) for weak supervision, and (III) a collection of unlabeled document sentences with which to build a training set. Ontologies instantiate labeling function templates that are applied to sentences to generate a label matrix. This matrix is used to train the label model that learns source accuracies and corrects for label noise to predict a consensus probability per word. Consensus labels are transformed into the probabilistic sequence label dataset that is used as training data for an end model (e.g., BioBERT). Alternatively, the label model can also be used as the final classifier.

Table 1 F1 scores for ontology and task-specific rule-based weak supervision.

Task	Ontologies (guidelines + UMLS + other)				+Task-specific rules				Hand-labeled	
	LFs	MV	LM	WS	LFs	MV	LM	WS	FS	SOTA
Chemical	22	79.8	88.0 ± 0.1 [†]	88.5 ± 0.2*	+9	81.1	89.2 ± 0.2 [†]	91.1 ± 0.1*	92.4 ± 0.2	93.5 ²⁴
Disease	16	74.7	78.9 ± 0.1[†]	78.3 ± 0.2*	+6	76.4	79.8 ± 0.3 [†]	79.9 ± 0.2	84.5 ± 0.2	87.2 ²⁴
Disorder	25	67.8	68.3 ± 0.3 [†]	69.1 ± 0.2*	+11	71.2	75.0 ± 0.2 [†]	76.3 ± 0.1*	79.6 ± 0.3	80.1 ⁶⁵
Drug	16	75.3	78.6 ± 0.1 [†]	79.2 ± 0.2*	+11	82.2	85.8 ± 0.4 [†]	88.3 ± 0.3*	93.2 ± 0.3	91.4 ⁶⁶
Negation	-	-	-	-	17	92.5	93.0 ± 0.0[†]	92.7 ± 0.6*	96.1 ± 0.2	-
DocTimeRel	-	-	-	-	27	67.8	69.2 ± 0.0 [†]	72.9 ± 0.5*	86.2 ± 0.1	83.4 ⁶⁷

Models are majority vote (MV); label model (LM); weakly supervised BioBERT (WS); our fully supervised BioBERT (FS); and published state-of-the-art (SOTA). LFs denote labeling function counts or total added task-specific rules. Bold indicates the best score for each approach and task. Scores are the mean and ± 1 SD of $n = 10$ random weight initializations. A two-sided Wilcoxon signed-rank test was used to compute statistical significance. *Denotes $P < 0.05$ for the difference between weakly supervised BioBERT (WS) and the label model (LM). For (chemical, disease, disorder, drug) exact P values for ontologies were (0.0039, 0.0020, 0.0020, 0.0020) and for task-specific rules (0.0020, 0.3223, 0.0020, 0.0020). For Negation $P = 0.0273$ and for DocTimeRel $P = 0.0020$. [†]Denotes $P < 0.05$ for the difference between the label model (LM) and majority vote (MV). Here, all task P values were 0.0020. *Mowery et al.⁶⁸ only reported accuracy for the negation task.

easily automated ontology-based weak supervision vs. more expensive, task-specific rules. Finally, we present a case study deploying Trove for COVID-19 symptom tagging and risk factor monitoring using a daily data feed of Stanford Health Care emergency department notes.

Weakly supervised learning is an umbrella term referring to methods for training classifiers using imperfect, indirect, or limited labeled data and includes techniques such as distant supervision^{26,27}, co-training²⁸, and others²⁹. Prior approaches for weakly supervised NER such as cotraining use a small set of labeled seed examples³⁰ that are iteratively expanded through bootstrapping or self-training³¹. Semi-supervised methods also use some amount of labeled training data and incorporate unlabeled data by imposing constraints on properties such as expected label distributions³². Distant supervision requires no labeled training data, but typically focuses on a single source for labels such as AutoNER²², which used phrase mining and a tailored dictionary of canonical entity names to construct a more precise labeler, rather than unifying labels assigned using heterogeneous sources of unknown quality. Crowdsourcing methods combine labels from multiple human annotators with unknown accuracy³³. However, compared to human labelers, programmatic label assignment has different correlation and scaling properties that create technical challenges when combining sources. Data programming^{11,17,34} formalizes theory for combining multiple label sources with different coverage and unknown accuracy, as well as correlation structure to correct for labeling errors.

In the setting of weakly supervised NER and sequence labeling, SwellShark²¹ uses a variant of data programming to train a generative model using labels from multiple dictionary and rule-based sources. However, this approach required task-specific preprocessing to identify candidate entities a priori to achieve competitive performance. Safranchik et al.²³ presented WISER, a linked hidden Markov model where weak supervision was defined separately over tags and tag transitions using linking rules derived from language models, ngram statistics, mined phrases, and custom heuristics to train a BiLSTM-CRF. SwellShark and WISER both focused on hand-coded, task-specific labeling function design.

Trove advances weakly supervised medical entity classification by (1) eliminating the requirement for identifying probable entity spans a priori by combining word-level weak supervision with contextualized word embeddings, (2) developing general-purpose, more easily automated ontology-based labeling functions that reduce the need for engineering hand-coded rules, (3) quantifying the relative contributions of sources of label assignment—such as pre-existing ontologies from the UMLS (low cost) and task-specific rule engineering (high cost)—to the achieved

performance for a task, and (4) evaluating Trove in a deployed medical setting, tagging symptoms and risk factors of COVID-19.

Results

Experiment overview. After quantifying the performance of ontology-driven weak supervision in all our tasks, we performed four experiments. First, we examined performance differences by label source ablations, which compared ontology-based labeling functions against those incorporating task-specific rules. Second, we compared Trove to existing weakly supervised tagging methods. Third, we examined learning source accuracies for UMLS terminologies. Finally, we report on a case study that used Trove to monitor emergency department notes for symptoms and risk factors associated with patients tested for COVID-19.

We evaluated four methods of combining labeling functions to train entity classifiers. (1) Majority vote (MV) is the majority class for each word predicted by all labeling functions. In cases of abstain or ties, predictions default to the majority class. (2) Label model (LM) is the default data programming model. Abstain and ties default to the majority class. (3) Weakly supervised (WS) is BioBERT trained on the probabilistic dataset generated by the label model. (4) Fully supervised (FS) is BioBERT trained on the original expert-labeled training set, tuned to match current published state-of-the-art performance, and using the validation set for early stopping.

For reference, we also included published F1 metrics for state-of-the-art (SOTA) supervised performance for each task, as determined to the best of our knowledge. Note that some published SOTA benchmarks (e.g., BC5CDR in Lee et al.²⁴) use both the hand-labeled train and validation sets for training, so they are not directly comparable to our experimental setup.

Performance of Trove in medical entity classification tasks.

Table 1 reports F1 performance for weak supervision using ontology-based labeling functions and those incorporating additional, task-specific rules. For NER tasks, adding task-specific rules performed within 1.3–4.9 F1 points (4.1%) of models trained on hand-labeled data and for span tasks within 3.4–13.3 F1 points. The total number of task-specific labeling functions used ranged from 9 to 27. For ontology-based supervision, the label model improved performance over MV by 4.1 F1 points on average, and BioBERT provided an additional average increase of 0.3 F1 points.

Labeling source ablations. For NER tasks, we examined five ablations, ordered by increasing the cost of labeling effort. (1) Guidelines, a dictionary of all positive and negative examples explicitly provided in annotation guidelines, including

dictionaries for punctuation, numbers, and English stopwords. (2) +UMLS, all terminologies available in the UMLS. (3) +Other, additional ontologies or existing dictionaries not included in the UMLS. (4) +Rules, task-specific rules, including regular expressions, small dictionaries, and other heuristics. (5) Hand-labeled, supervised learning using the expert-labeled training split.

Tiers 1–4 are additive and include all prior levels. We initialized labeling function templates as follows:

For ontology-based labeling functions, we used the UMLS Semantic Network and the corresponding Semantic Groups as our entity categories and defined a mapping of semantic types (STYs) to target class labels $y \in \{-1, 0, 1\}$. Non-UMLS ontologies that did not provide semantic-type assignments (e.g., ChEBI) were mapped to a single class label. All UMLS terminologies v were ranked by term coverage on the unlabeled training set, defined as each term's document frequency summed by terminology, and the top s terminologies were used to initialize templates, where s was tuned with a validation set. The remaining (v_{s+1}, \dots, v_{92}) UMLS terminologies were merged into a single labeling function to ensure that all terms in the UMLS were included. UMLS synsets were constructed using concept unique identifiers (CUIs) and templates were initialized with the union of all terminologies and fixed across all NER tasks.

For task-specific labeling functions, we evaluated our ability to supplement ontology-based supervision with hand-coded labeling functions and estimated the relative performance contribution of adding these task-specific rules. All training set documents were preprocessed to tag entities using the ontology-based labeling functions outlined above and indexed to support search queries for efficient data exploration. The design of task-specific labeling functions is a mix of data exploration, i.e., looking at entities identified by ontology labeling functions to identify errors, and similarity search to identify common, out-of-ontology concept patterns. Only the training set was examined during this process and the test set was held out during all labeling function development and model tuning.

For NER, we used two rule types to label concepts: (1) pattern matching via regular expressions and small dictionaries of related terms (e.g., illegal drugs), and (2) bigram word co-occurrence graphs from ontologies to support fuzzy span matching. Pattern matching comprised the majority of our task-specific labeling functions. While task-specific labeling functions codify generalized patterns not captured by ontologies, we also note that a number of our task-specific labeling functions were necessary due to the idiosyncratic nature of ground-truth labels in benchmark tasks. For example, in the i2b2/n2c2 drug tagging task, annotation guidelines included more complex, conditional entity definitions, such as not labeling negated or historical drug mentions. We incorporated these guidelines using the Negation and DocTimeRel labeling functions described below. See Supplementary Fig. 1 and Supplementary Note for a more detailed example of designing task-specific labeling functions.

For span tasks, which classify Negation and DocTimeRel for preidentified entities, we do not use ontology-based labeling functions directly for supervision. Instead, ontology-tagged entities were used to guide the development of labeling functions that search left- and right-context windows around a target entity for cue phrases. Designing search patterns for left- and right-context windows is the same strategy used by NegEx/ConText^{12,35} to assign negation and temporal status. For Negation, we built on NegEx by adding additional patterns found via exploration of the training documents. For DocTimeRel, we used a heuristic based on the nearest explicit datetime mention (in the token distance) to an event mention³⁶. Additional contextual pattern matching rules were added to detect other cues of event temporality, e.g., using section headers such as past medical history to identify events occurring before the note creation time.

Figure 2 reports F1 scores across all ablation tiers. In all settings, the weakly supervised BioBERT models outperformed MV. Gains of 8.0–34.7 F1 points are seen in the guideline-only tier and 1.3–8.2 points in other tiers. Incorporating source




Fig. 2 Ablation study of F1 performance by labeling source. Majority vote (MV) vs. weakly supervised BioBERT (WS) vs. fully supervised (FS) for all labeling source ablations showing the absolute F1 score for all labeling tiers. The colored region of each bar indicates MV performance and the white regions denote performance improvements of WS over MV. The mean performance of FS is indicated by the green lines and square points. WS and FS consist of $n = 10$ experiment replicates using different random initialization seeds, presented as the mean with error bars \pm SD. MV is deterministic and does not include replicates.

accuracies into BioBERT training provided significant benefits when combining high-precision sources with low-precision/high-recall sources. In the case of chemical tagging with MV, the UMLS tier (red) outperformed UMLS+Other (orange) by 1.8 F1 points (81.6 vs. 79.8). This was due to adding the ChEBI ontology that increased recall but only had 65% word-level precision. Majority vote cannot learn or utilize this information, so naively adding ChEBI labels hurts performance. However, the label model learned ChEBI's accuracy to take advantage of the noisier, but higher-coverage signal; thus, the WS UMLS + Other (orange + white) outperformed UMLS ((red + white)) by 2.5 F1 points (88.0 vs. 85.5). See Supplementary Tables 1–4 for complete performance metrics across all ablation tiers.

Comparing Trove with existing weakly supervised methods.

We compared Trove to three existing weakly supervised methods for NER and sequence labeling: SwellShark²¹, AutoNER²², and WISER²³. We compared performance on BC5CDR (the combination of disease and chemical tasks) against all methods and on the i2b2/n2c2 drug task for SwellShark. All performance numbers are for models trained on the original training set split, with the exception of SwellShark that is trained on an additional 25,000 weakly labeled documents. All weakly supervised methods use the labeling functions, preprocessing, and dictionary curation methods as described in the original manuscripts. Table 2 compares Trove with these existing weakly supervised methods. Our ontology-based approach outperformed AutoNER by 1.7 F1 points. For models incorporating task-specific rules, we outperformed the best weakly supervised model SwellShark by 1.9 F1 points. SwellShark reported F1 scores on the i2b2/n2c2 drug task of 78.3 for dictionaries and 83.4 for task-specific rules. Our best models achieved 79.2 and 88.4 F1, respectively.

UMLS terminologies as plug-and-play weak supervision. Bio-medical annotators such as NCBO BioPortal require selecting a set of target ontologies/terminologies to use for labeling. Since Trove is capable of automatically combining noisy terminologies, given a shared semantic-type definition, we tested the ability to avoid selecting specific UMLS terminologies for use as supervision sources. This is challenging because estimating accuracies with the label model requires observing agreement and disagreement among multiple label sources; however, it is non-obvious how to partition the UMLS, which contains many terminologies, into labeling functions. The naive extremes are to either create a single labeling function from the union of all terminologies or include all terminologies as individual labeling functions.

To explore how partitioning choices impact label model performance, we held all non-UMLS labeling functions fixed

across all ablation tiers and computed performance across $s = (1, \dots, 92)$ partitions of the UMLS by terminology. All scores were normalized to the best global majority vote score per tier, selected using the best s choice evaluated on the validation set, to assess the impact of correcting for label noise.

Figure 3 shows the impact of partitioning the UMLS into s different labeling functions. Modeling source accuracy consistently outperformed MV across all tiers, in some cases by 2–8 F1 points. The best performing partition size s ranged from 1 to 10 by task. The naive baseline approaches—collapsing the UMLS into a single labeling function or treating all terminologies as individual labeling functions—generally did not perform best overall.

Case study in rapidly building clinical classifiers. We deployed Trove to monitor emergency departments for patients undergoing COVID-19 testing, analyzing clinical notes for presenting symptoms/disorders and risk factors³⁷. This required identifying disorders and defining a novel classification task for exposure to a confirmed COVID-19-positive individual, a risk factor informing patient contact tracing. The dataset consisted of daily dumps of emergency department notes from Stanford Health Care (SHC), beginning in March 2020. Our study was approved by the Stanford University Administrative Panel on Human Subjects Research, protocol #24883, and included a waiver of consent. All included patients from SHC signed a privacy notice, which informs them that their records may be used for research purposes given approval by the IRB, with study procedures in place to protect patient confidentiality.

We manually annotated a gold test set of 20 notes for all mentions of disorders and 776 notes for mentions of positive COVID exposure. Two clinical experts generated gold annotations that were adjudicated for disagreements by authors AC and JAF. As a baseline for disorder tagging, we used the fully supervised ShARE/CLEF disorder tagger. This reflects a readily available, but out-of-distribution training set (MIMIC-II³⁸ vs. SHC). We used the same disorder labeling function set as our prior experiments, adding one additional dictionary of COVID terms³⁹. BioBERT was trained using 2482 weakly labeled documents. Custom labeling functions were written for the exposure task and models were trained on 14 k sentences.

Table 3 contains our COVID case study results. The label model provided up to 5.2 F1 points improvement over majority vote and performed best overall for disorder tagging. Our best weakly supervised model outperformed the disorder tagger trained on hand-labeled MIMIC-II data by 2.3 F1 points. For exposure classification, the label model provided no benefit, but the weakly supervised end model provided a 6.9% improvement (+5.2 F1 points) over the rules alone.

Table 2 Comparison of Trove against existing weakly supervised NER methods.

Supervision method	Label source	Number of train docs	End model	P	R	F1
Fully supervised	Hand-labeled	500	BioBERT	87.6	89.3	88.7
Fully supervised	Hand-labeled	500	BiLSTM-CRF	87.2	87.9	87.5
SwellShark	Dictionaries	25,500	BiLSTM-CRF	<u>84.6</u>	74.1	79.0
AutoNER	Dictionaries	500	BiLSTM-CRF	<u>83.2</u>	81.1	82.1
Ours (Trove + Snorkel)	Dictionaries	500	BioBERT	81.6	<u>86.1</u>	<u>83.7</u>
SwellShark	Custom rules	25,500	BiLSTM-CRF	86.1	82.4	84.2
WISER	Custom rules	500	BiLSTM-CRF	82.7	83.3	83.0
Ours (Trove + Snorkel)	Custom rules	500	BioBERT	85.5	86.8	86.1

Precision (P), recall (R), and F1 scores for the BC5CDR task. Underlined numbers indicate the best weakly supervised score using only dictionaries/ontologies, and bold indicates the best score using custom rules. For this task, ontology-based supervision alone outperformed existing weakly supervised methods except for SwellShark which required custom rules and candidate generation. Incorporating task-specific rules into Trove further improved performance.

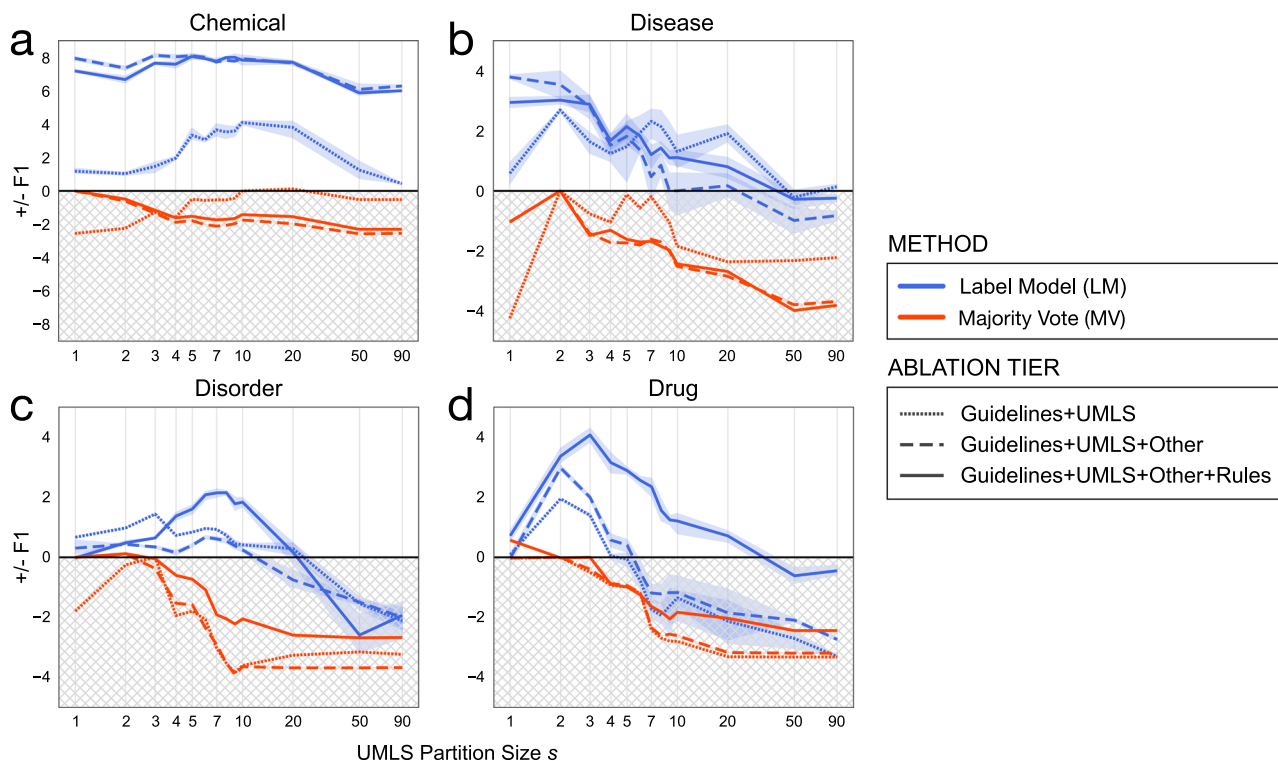


Fig. 3 The relationship between the number of UMLS partitions and the learned accuracies of label sources. **a** BC5CDR chemical entities. **b** BC5CDR disease entities. **c** ShARe/CLEF 2014 disorder entities. **d** i2b2/n2c2 2009 drug entities. The UMLS is partitioned into *s* terminologies (*x* axis, log-scale) ordered by term coverage on the unlabeled training set. Red (MV) and blue (LM) lines are the mean difference in F1 performance (*y* axis) of *n* = 5 random weight initializations. Error bars are represented using the solid colored line to denote the mean value of data points and the shaded regions corresponding to ± SD. The gray region indicates performance worse than the best possible MV, discovered via the validation set. Across virtually all partitioning choices, modeling source accuracies outperformed MV, with *k* = 1–10 performing best overall.

Table 3 COVID-19 presenting symptoms/disorders and risk factors evaluated on Stanford Health Care emergency department notes.

Supervision	Task	MV			LM			WS			FS		
		P	R	F1	P	R	F1	P	R	F1	P	R	F1
Hand-labeled	Disorder	-	-	-	69.3	71.7	70.5	67.1	72.3	69.6	68.0	74.5	71.1
Ontologies	Disorder	64.4	66.4	65.3	73.0	73.9	73.4	70.5	74.8	72.6	-	-	-
+Task-specific	Disorder	69.1	70.4	69.8	73.0	73.9	73.4	70.5	74.8	72.6	-	-	-
Task-specific	Exposure	82.6	69.1	75.2	82.6	69.1	75.2	<u>87.2</u>	<u>74.5</u>	<u>80.4</u>	-	-	-

Bold and underlined scores indicate the best score in symptom/disorder tagging and COVID exposure classification respectively. Ontology-based weak supervision performed almost as well as the out-of-distribution, hand-labeled MIMIC-II data used for FS. Adding task-specific rules, even though they were developed without seeing Stanford data, outperformed the hand-labeled FS model by 2.3 F1 points.

Discussion

Our experiments demonstrate the effectiveness of using weakly supervised methods to train entity classifiers using off-the-shelf ontologies and without requiring hand-labeled training data. Medical ontologies are freely available sources of weak supervision for NLP applications⁴⁰, and in several NER tasks, our ontology-only weakly supervised models matched or outperformed more complex weak supervision methods in the literature. Our work also highlights how domain-aware language models, such as BioBERT, can be combined with weak supervision to build low-cost and highly performant medical NLP classifiers.

Rule-based approaches are common tools in scientific literature analysis and clinical text processing⁴¹. Our results suggest that engineering task-specific rules in addition to labels provided by ontologies provides strong performance for several NER tasks—in

some cases approaching the performance of systems built using hand-labeled data. We further demonstrated how leveraging the structure inherent in knowledge bases such as the UMLS to estimate source accuracies and correct for label noise provides substantial performance benefits. We find that the classification performance of the label model alone is strong, with BioBERT providing modest gains of 1.0 F1 points on average. Since the label model is orders of magnitude more computationally efficient to train than BERT-based models, in many settings (e.g., limited access to high-end GPU hardware), the label model alone may suffice.

Our tasks reflect a wide range of difficulties. Clinical tasks required more task-specific rules to address the increased complexity of entity definitions and other nongrammatical, sub-language phenomena⁴². Here custom rules improved clinical tasks an average of 8.1 F1 points vs. 2.1 points for scientific

Table 4 Dataset summary statistics.

Task	Domain	Name	Type	k	Documents	Entities
Disease	Literature	BC5CDR ⁶⁹	NER	2	500/500/500	4182/4244/4424
Chemical	Literature	BC5CDR ⁶⁹	NER	2	500/500/500	5203/5347/5385
Disorder	Clinical	ShARe/CLEF 2014 ⁶⁸	NER	2	166/133/133	5619/4449/7367
Drug	Clinical	i2b2/n2c2 2009 ⁷⁰	NER	2	100/75/75	3157/2504/2819
Negation	Clinical	ShARe/CLEF 2014 ⁶⁸	Span	2	166/133/133	5619/4449/7367
DocTimeRel	Clinical	THYME 2016 ⁷¹	Span	4	293/147/151	38937/20974/18990

There are (k) classes per task. The (Documents) and (Entities) columns indicate counts for train/validation/test splits.

literature. Moreover, adding non-UMLS ontologies to PubMed tasks consistently improved overall performance while providing little-to-no benefit for our clinical tasks. Annotation guidelines for our clinical tasks also increased complexity. The i2b2/n2c2 drug task combines several underlying classification problems (e.g., filtering out negated medications, patient allergies, and historical medications) into a single tagging formulation. This extends beyond entity typing and requires a more complex cue-driven rule design.

Manually labeling training data is time-consuming and expensive, creating barriers to using machine learning for new medical classification tasks. Sometimes, there is a critical need to rapidly analyze both scientific literature and unstructured electronic health record data—as in the case of the COVID-19 pandemic when we need to understand the full repertoire of symptoms, outcomes, and risk factors at short notice^{37,43,44}. However, sharing patient notes and constructing labeled training sets presents logistical challenges, both in terms of patient privacy and in developing infrastructure to aggregate patient records⁴⁵. In contrast, labeling functions can be easily shared, edited, and applied to data across sites in a privacy-preserving manner to rapidly construct classifiers for symptom tagging and risk factor monitoring.

This work has several limitations. Our task-specific labeling functions were not exhaustive and only reflect low-cost rules easily generated by domain experts. Additional rule development could lead to improved performance. In addition, we did not explore data augmentation or multitask learning in the BioBERT model, which may further mitigate the need to engineer task-specific rules. There is considerable prior work developing machine-learning models for tagging disease, drug, and chemical entities that could be incorporated as labeling functions. However, our goal was to explore performance trade-offs in settings where existing machine-learning models are not available. Our framework leverages the wide range of medical ontologies available for English language settings, which provides considerable advantages for weakly supervised methods. Additional work is needed to characterize the extent to which the framework can benefit tasks in non-English settings. Combining labels from multiple ontology sources violates an independence assumption of data programming as used in this work, because for any pair of source ontologies, we may have correlated noise. This restriction applies to all label sources but is more prevalent in cases with extremely similar label sources, as can occur with ontologies. In our experiments, for a small number of sources, the impact was minor; however, performance tended to decrease after including more than 20 ontologies. Additional research into unsupervised methods for structure learning^{46,47}, i.e., learning dependencies among sources from unlabeled data, could further improve performance or mitigate the need to limit the number of included ontologies.

Identifying named entities and attributes, such as negation, are critical tasks in medical natural language processing. Manually

labeling training data for these tasks is time consuming and expensive, creating a barrier to building classifiers for new tasks. The Trove framework provides ontology-driven weak supervision for medical entity classification and achieves state-of-the-art weakly supervised performance in the NER tasks of recognizing chemicals, diseases, and drugs. We further establish new weakly supervised baselines for disorder tagging and classifying the temporal order of an event entity relative to its document timestamp. The weakly supervised NER classifiers perform within 1.3–4.9 F1 points of classifiers trained with hand-labeled data. Modeling the accuracies of individual ontologies and rules to correct for label noise improved performance in all of our entity classification tasks. Combining pretrained language models such as BioBERT with weak supervision results in an additional improvement in most tasks.

The Trove framework demonstrates how classifiers for a wide range of medical NLP tasks can be quickly constructed by leveraging medical ontologies and weak supervision without requiring manually labeled training data. Weakly supervised learning provides a mechanism for combining the generalization capabilities of state-of-the-art machine learning with the flexibility and inspectability of rule-based approaches.

Methods

Datasets and tasks. We analyze two categories of medical tasks using six datasets: (1) NER and (2) span classification where entities are identified a priori and classified for cue-driven attributes such as negation or document relative time, i.e., the order of an event entity relative to the parent document's timestamp. Both categories of tasks are formalized as token classification problems, either tagging all words in a sequence (NER) or just the head words for an entity set (span classification). Table 4 contains summary statistics for all six datasets. All documents were preprocessed using a spaCy⁴⁸ pipeline optimized for biomedical tokenization and sentence boundary detection¹⁹.

Our COVID-19 case study used a daily feed of emergency department notes from Stanford Health Care (SHC), beginning in March 2020. Our study was approved by the Stanford University Administrative Panel on Human Subjects Research, protocol #24883 and included a waiver of consent. All included patients from SHC signed a privacy notice which informs them that their records may be used for research purposes given approval by the IRB, with study procedures in place to protect patient confidentiality.

We used 99 label sources covering a broad range of medical ontologies. We used the 2018AA release of the UMLS Metathesaurus, removing non-English and zoonotic source terminologies, as well as sources containing fewer than 500 terms, resulting in 92 sources. Additional sources included the 2019 SPECIALIST abbreviations⁴⁹, Disease Ontology⁵⁰, Chemical Entities of Biological Interest (ChEBI)⁵¹, Comparative Toxicogenomics Database (CTD)⁵², the seed vocabulary used in AutoNER²², ADAM abbreviations database⁵³, and word sense abbreviation dictionaries used by the clinical abbreviation system CARD⁵⁴.

We applied minimal preprocessing to all source ontologies, filtering out English stopwords and numbers, applying a letter case normalization heuristic to preserve abbreviations, and removing all single-character terms. We did not incorporate UMLS term-type information, such as filtering out terms explicitly denoted as suppressible within a terminology since this information is not typically available in non-UMLS ontologies. Our overall goal was to impose as few assumptions as possible when importing terminologies, evaluating their ability to function as plug-and-play sources for weak supervision.

Formulation of the labeling problem. We assume a sequence-labeling problem formulation, where we are given a dataset $D = \{X_i\}_{i=1}^N$ of N sequences

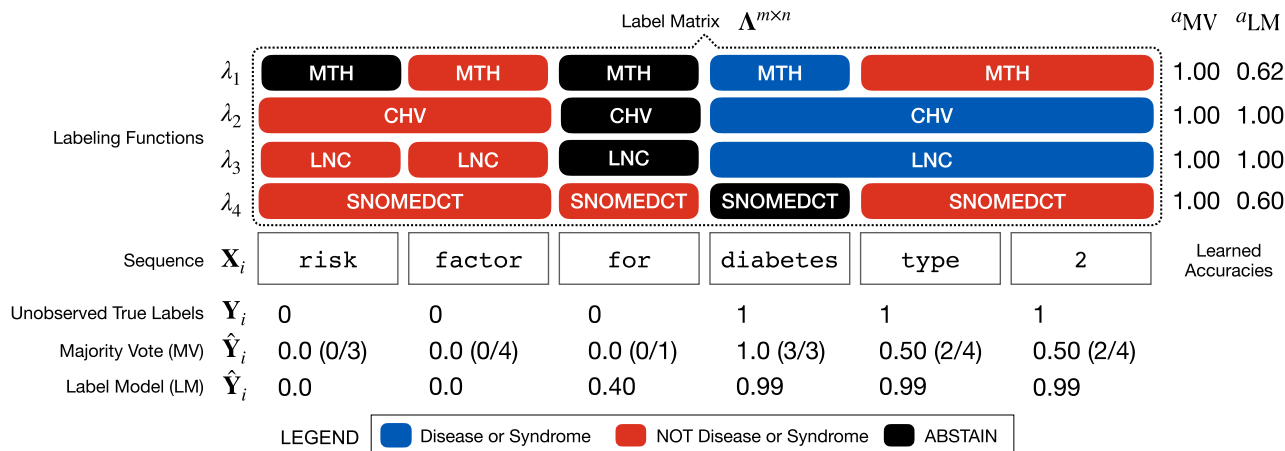


Fig. 4 An example of combining ontology-based labeling functions. Here four ontology labeling functions (MTH, CHV, LNC, SNOMEDCT) are used to label a sequence of words X_i , containing the entity diabetes type 2. Majority vote estimates Y_i as a word-level sum of positive class labels, weighing each equally (a_{MV}). The label model learns a latent class-conditional accuracy (a_{LM}) for each ontology, which is used to reweight labels to generate a more accurate consensus prediction of Y_i .

$X_i = (x_{i,1}, \dots, x_{i,t})$ consisting of words x from a fixed vocabulary. Each sequence is mapped to a corresponding sequence of latent class variables $Y_i = (y_{i,1}, \dots, y_{i,t})$, where $y \in \{0, \dots, k\}$ for k -tag classes. Since Y is not observable, our primary technical challenge is estimating Y from multiple, potentially conflicting label sources of unknown quality to construct a probabilistically labeled dataset $\hat{D} = \{X_i, \hat{Y}_i\}_{i=1}^N$. This dataset can then be used for training classification models such as deep neural networks. Such a labeling regimen is typically low-cost, but less accurate than the hand-curated labels used in traditional supervised learning; hence, this paradigm is referred to as weakly supervised learning.

Unifying and denoising sources with a label model. When using biomedical annotators such as MetaMap or NCBO BioPortal, users specify a target set of entity classes and a set of terminology sources with which to generate labeled concepts. Consider the example outlined in Fig. 4, where we want to train an entity tagger for disease names using labels generated from four terminologies. Here, we are interested in generating a consensus set of entities using each terminology’s labeled output. A straightforward unification method is majority vote

$$\hat{y} = \arg \max_{y \in \{1, \dots, k\}} \sum_{i=1}^m \mathbb{1}_k(\lambda_i(x) = y) \quad (1)$$

where our m terminologies are represented as individual labeling functions λ_i . Labeling functions encode an underlying heuristic such as matching strings against a dictionary and given an input instance (e.g., a document or entity span) assign a label in the domain $\{-1, 0, \dots, k\}$ where -1 denotes abstain, i.e., not assigning any class label. Majority vote simply takes the mode of all labeling function outputs for each word, emitting the majority class in the case of ties or abstains.

Majority vote weights source equally when combining labels, an assumption that does not hold in practice, which introduces noise into the labeling process. Sources have unknown, task-dependent accuracies and often make systematic labeling errors. Failing to account for these accuracies can negatively impact classification performance. To correct for such label noise, we use data programming³⁴ to estimate accuracies of each source and ensemble the sources via a label model that assigns a consensus probabilistic label per word.

To learn the label model, m -label sources are parameterized as labeling functions $\lambda_1, \dots, \lambda_m$. The vector of m -labeling functions applied to n instances forms the label matrix $\Lambda \in \{-1, 0, \dots, k\}^{m \times n}$. A key finding of data programming is that we can use Λ to recover the latent class-conditional accuracy of each label source without ground-truth labels by observing the rates of agreement and disagreement across all pairs of labeling functions λ_i, λ_j ³⁴. This leverages the fact that while the accuracy $a_i = \mathbb{E}[\lambda_i Y]$ (the expectation of the labeling function output λ_i , multiplied by the true label) is not directly observable, the product of $a_i a_j = \mathbb{E}[\lambda_i Y \lambda_j Y] = \mathbb{E}[\lambda_i Y] \mathbb{E}[\lambda_j Y]$ is the rate at which labeling functions vote together, which is observable via Λ . Assuming independent noise among labeling functions, accuracies are then recoverable up to a sign by solving accuracies for disjoint sets of triplets. We refer readers to Ratner et al.¹⁷ for more details.

We use the weak supervision framework Snorkel¹¹ to train a probabilistic label model that captures the relationship between the true label and label sources $P(Y, \Lambda)$. Here, the training input is the label matrix Λ , generated by applying labeling functions $\lambda_1, \dots, \lambda_m$ to the unlabeled dataset D . Formally, $P(Y, \Lambda)$ can be encoded as a factor graph-based model with m -accuracy factors between $\lambda_1, \dots, \lambda_m$ and our

true (unobserved) label y (Fig. 1, step 3).

$$\theta_j^{Acc}(\Lambda_i, y_i) := y_i \Lambda_{ij} \quad (2)$$

$$p_\theta(Y, \Lambda) \propto \exp\left(\sum_{i=1}^m \sum_{j=1}^n \theta_j^{Acc} \phi_j^{Acc}(\Lambda_i, y_i)\right) \quad (3)$$

Snorkel implements a matrix completion formulation of data programming, which enables faster estimation of model parameters θ using stochastic gradient descent rather than relying on Gibbs sampling-based approaches¹⁷. The label model estimates $P(Y|\Lambda)$ to provide denoised consensus label predictions \tilde{Y} and generates our probabilistically labeled dataset \tilde{D} .

Figure 4 shows how data programming provides a principled way to synthesize a label when there is disagreement across label sources about what constitutes an entity span. The disease mentioning diabetes type 2 is not found in Metathesaurus Names (MTH) or SNOMED Clinical Terms (SNOMEDCT), which leads to disagreement and label errors. Using a majority vote of labeling functions misses the complete entity span, while the label model learns to account for systematic errors made by each ontology to generate a more accurate consensus label prediction.

Labeling function templates. In this work, a labeling function λ_j accepts an unlabeled sequence X_i as input and emits a vector of predicted labels

$\tilde{Y}_{i,j} = (\tilde{y}_{j,1}, \dots, \tilde{y}_{j,t})$, i.e., a label $\tilde{y}_j \in \{-1, 0, \dots, k\}$ for each word in X_i . A typical labeling function serves as a wrapper for an underlying, potentially task-specific labeling heuristic such as pattern matching with a regular expression or a more complex rule system. Since these labeling functions are not easily automated and require hand coding, we refer to them as task-specific labeling functions. These are analogous to the rule-based approaches used in 48% of recent medical concept recognition publications⁴¹.

In contrast, medical ontologies can be automatically transformed into labeling functions with little-to-no custom coding by defining reusable labeling function templates. Templates only require specifying a set of target entity categories and providing a collection of terminologies mapped to those categories. These categories are easily derived from knowledge bases such as the UMLS Metathesaurus (where the UMLS Semantic Network⁵⁵ provides consistent categorization of UMLS concepts) or other domain-specific taxonomies. In this work, we use UMLS Semantic Groups⁵⁶ (mappings of semantic types into simpler, nonhierarchical categories such as disorders) as the basis for our concept categories.

We explore two types of ontology-based labeling functions, which leverage knowledge codified in medical ontologies for term semantic types and synonymy.

Semantic type labeling functions require a set of terms (single or multiword entities) $t \in T$ mapped to semantic types, where a term may be mapped to multiple-entity classes. This mapping is converted to a k -dimensional probability vector where k is the number of entity classes $t_i \rightarrow [p_1, \dots, p_k]$. Given input sequence X_i , use string matching to find all longest-term matches (in token length) and assign each match to its most probable entity class $\tilde{y} = \max(t_i)$, abstaining on ties. Using the longest match is a heuristic that helps disambiguate nested terms (lung as anatomy vs. lung cancer as a disease). Matching optionally includes a set of slot-filled patterns to capture simple compositional mentions (e.g., $\{*\} (\{*\}) \rightarrow$ Tylenol (Acetaminophen)).

Synonym (synset) labeling functions require synsets (collections of synonymous terms) $\{\hat{t}_1, \dots, \hat{t}_n\} \in \hat{T}$ and term T mapped to semantic types. Given input sequence X_i and its parent context (e.g., document) search for > 1 unique synonym matches from a target synset and label all matches $\hat{y} = \max(\hat{t}_i)$. This is useful for disambiguating abbreviations (e.g., Duchenne muscular dystrophy \rightarrow DMD), where a long-form of an abbreviated term appears elsewhere in a document. Matches can be unconstrained, e.g., any tuple found anywhere in a context, or subject to matching rules, e.g., using Schwartz–Hearst abbreviation disambiguation⁵⁷ to identify out-of-dictionary abbreviations.

Training the BioBERT end model. The output of the label model is a set of probabilistically labeled words, which we transform back into sequences $\hat{D} = \{X_i, \hat{Y}_i\}_{i=1}^N$. While probabilistic labels may be used directly for classification, this suffers from a key limitation: the label model cannot generalize beyond the direct output of labeling functions. Rules alone can miss common error cases such as out-of-dictionary synonyms or misspellings. Therefore, to improve coverage, we train a discriminative end model, in this case a deep neural network, to transform the output of labeling functions into learned feature representations. Doing so leverages the inductive bias of pretrained language models⁵⁸ and provides additional opportunities for injecting domain knowledge via data augmentation⁵⁹ and multitask learning⁶⁰ to improve classification performance.

We use the transformer-based BioBERT²⁴, a language model fine-tuned on the biomedical text. We also evaluated ClinicalBERT⁶¹ for clinical tasks, and found its performance to be the same as BioBERT. BioBERT is trained as a token-level classifier with a max sequence length of 512 tokens. We follow Devlin et al.⁵⁸ for sequence-labeling formulation, using the last BERT layer of each word's head wordpiece token as the contextualized embedding. Since sequence labels may be incomplete (i.e., cases where all labeling functions abstain on a word), we mask all abstained tokens when computing the loss during training. We modified BioBERT to support a noise-aware binary cross-entropy loss function³⁴ that minimizes the expected value with respect to \hat{Y} to take advantage of the more informative probabilistic labels.

$$\hat{w} = \underset{w}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{y \sim \hat{Y}} [L(w, x_i, \hat{y})] \quad (4)$$

Hyperparameter tuning for the label and end models. All models were trained using weakly labeled versions of the original training splits, i.e., no hand-labeled instances. We used a hand-labeled validation and test set for hyperparameter tuning and model evaluation, respectively. Result metrics are reported using the test set. The label model was tuned for learning rate, training epochs, L2 regularization, and a uniform accuracy prior used to initialize labeling function accuracies. BioBERT weights were fine-tuned, and end models were tuned for learning rate and training epochs. We used a linear decay learning rate schedule with a 10% warmup period. See Supplementary Tables 5 and 6 for hyperparameter grids.

Metrics. We report precision, recall, and F1 score for all tasks. DocTimeRela is reported using microaveraging. NER metrics are computed using exact span matching⁶². Each NER task is trained separately as a binary classifier using IO (inside, outside) tagging to simplify labeling function design, with predicted tags converted to BIO (beginning, inside, outside) to properly count errors detecting head words. Span task metrics are calculated assuming access to gold test set spans, as per the evaluation protocol of the original challenges. Label model and BioBERT scores are reported as the mean and standard deviation of 10 runs with different random seeds. A two-sided Wilcoxon signed-rank test with an alpha level of 0.05 was used to calculate statistical significance.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

All primary data that support the findings of this study are available via public benchmark datasets (BC5CDR, <https://biocreative.bioinformatics.udel.edu/tasks/biocreative-v/track-3-cdr/>) or are otherwise available per data use agreements with the respective data owners (ShARe/CLEF 2014, <https://physionet.org/content/shareclefhealth2014task2/1.0/>); THYME, <https://healthnp.hms.harvard.edu/center/pages/data-sets.html>; i2b2/n2c2 2009, <https://portal.dbmi.hms.harvard.edu/projects/n2c2-nlp/>). The data that support the findings of the clinical case study are available on request from the corresponding author J.A.F. These data are not publicly available because they contain information that could compromise patient privacy. Trove requires access to the UMLS, which is available by license from the National Library of Medicine, Department of Health and Human Services, <https://www.nlm.nih.gov/research/umls/index.html>. Open source ontologies used in this study are available at SPECIALIST Lexicon, <https://lsg3.nlm.nih.gov/LexSysGroup/Summary/lexicon.html>; Disease Ontology, <https://bioportal.bioontology.org/ontologies/DOID>; Chemical Entities of Biological Interest (ChEBI), <ftp://ftp.ebi.ac.uk/pub/databases/chebi/>; Comparative

Toxicogenomics Database (CTD), <http://ctdbase.org>; AutoNER core dictionary, https://github.com/shangjingbo1226/AutoNER/blob/master/data/BC5CDR/dict_core.txt; ADAM abbreviations database, http://arrowsmith.psych.uic.edu/arrowsmith_uic/adam.html; and the Clinical Abbreviation Recognition and Disambiguation (CARD) framework, <https://sbmi.uth.edu/ccb/resources/abbreviation.htm>.

Code availability

Trove is written in Python v3.6, spaCy 2.3.4 was used for NLP preprocessing, and Snorkel v0.9.5 was used for training the label model. BioBERT-Base v1.1, Transformers v2.8⁶³, and PyTorch v1.1.0 were used to train all discriminative models. Trove is open-source software and publicly available at <https://github.com/som-shahlab/trove>; <https://doi.org/10.5281/zenodo.4497214>.

Received: 7 August 2020; Accepted: 26 February 2021;

Published online: 01 April 2021

References

- Ravi, D. et al. Deep learning for health informatics. *IEEE J. Biomed. Health Informat.* **21**, 4–21 (2017).
- Esteva, A. et al. A guide to deep learning in healthcare. *Nat. Med.* **25**, 24–29 (2019).
- Wang, L. L. et al. COVID-19: The COVID-19 open research dataset. In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020* (eds Karin Verspoor, Kevin Bretonnel Cohen, Mark Dredze, Emilio Ferrara, Jonathan May, Robert Munro, Cecile Paris & Byron Wallace) (Association for Computational Linguistics, Online, 2020) <https://www.aclweb.org/anthology/2020.nlp-covid19-acl.1>.
- Kuleshov, V. et al. A machine-compiled database of genome-wide association studies. *Nat. Commun.* **10**, 3341 (2019).
- Fries, J. A. et al. Weakly supervised classification of aortic valve malformations using unlabeled cardiac MRI sequences. *Nat. Commun.* **10**, 3111 (2019).
- Khattar, S. et al. Multi-frame weak supervision to label wearable sensor data. in *Proceedings of the Time Series Workshop at ICML 2019* (eds Vitaly Kuznetsov, Cheng Tang, Yuyang Wang, Scott Yang & Rose Yu) (2019) <http://roseyu.com/time-series-workshop/>.
- Varma, P. et al. Multi-resolution weak supervision for sequential data. in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8–14 December 2019, Vancouver, BC, Canada.* (ed. Wallach, H. M. et al.) 192–203 (Neural Information Processing Systems, 2019).
- Dunnmon, J. A. et al. Cross-modal data programming enables rapid medical machine learning. *Patterns* **1**, 100019 (2020).
- Bodenreider, O. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic Acids Res.* **32**, D267–70 (2004).
- Jonquet, C., Shah, N. H. & Musen, M. A. The open biomedical annotator. *Summit Transl. Bioinform.* **2009**, 56–60 (2009).
- Ratner, A. et al. Snorkel: rapid training data creation with weak supervision. *Proc. VLDB Endowment* **11**, 269–282 (2017).
- Chapman, W. W., Bridewell, W., Hanbury, P., Cooper, G. F. & Buchanan, B. G. A simple algorithm for identifying negated findings and diseases in discharge summaries. *J. Biomed. Inform.* **34**, 301–310 (2001).
- Peng, Y. et al. NegBio: a high-performance tool for negation and uncertainty detection in radiology reports. *AMIA Jt Summits Transl Sci Proc* **2017**, 188–196 (2018).
- Wang, X. et al. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*, 3462–3471 <https://doi.org/10.1109/CVPR.2017.369> (IEEE Computer Society, 2017).
- Rajpurkar, P. et al. Deep learning for chest radiograph diagnosis: a retrospective comparison of the CheXNeXt algorithm to practicing radiologists. *PLoS Med.* **15**, e1002686 (2018).
- Draeos, R. L. et al. Machine-learning-based multiple abnormality prediction with large-scale chest computed tomography volumes. *Med. Image Anal.* **67**, 101857 (2020).
- Ratner, A. et al. Training complex models with multi-task weak supervision. *Proc. Conf. AAAI Artif. Intell.* **33**, 4763–4771 (2019).
- Wang, Y. et al. A clinical text classification paradigm using weak supervision and deep representation. *BMC Med. Inform. Decis. Mak.* **19**, 1 (2019).
- Callahan, A. et al. Medical device surveillance with electronic health records. *npj Digit Med.* **2**, 94 (2019).
- Peterson, K. J., Jiang, G. & Liu, H. A corpus-driven standardization framework for encoding clinical problems with HL7 FHIR. *J. Biomed. Inform.* **110**, 103541 (2020).

21. Fries, J., Wu, S., Ratner, A. & Ré, C. SwellShark: a generative model for biomedical named entity recognition without labeled data. Preprint at <https://arxiv.org/abs/1704.06360> (2017).
22. Shang, J. et al. Learning named entity tagger using domain-specific dictionary. in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, (eds Jingbo Shang, Liyuan Liu, Xiaotao Gu, Xiang Ren, Teng Ren & Jiawei Han) 2054–2064 (Association for Computational Linguistics, 2018).
23. Safranchik, E., Luo, S. & Bach, S. H. Weakly supervised sequence tagging from noisy rules. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020*, 5570–5578 (AAAI Press, 2020). <https://aaai.org/ojs/index.php/AAAI/article/view/6009>.
24. Lee, J. et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **36**, 1234–1240 (2019).
25. Aronson, A. R. & Lang, F.-M. An overview of metapmap: historical perspective and recent advances. *J. Am. Med. Inform. Assoc.* **17**, 229–236 (2010).
26. Craven, M. & Kumlien, J. Constructing biological knowledge bases by extracting information from text sources. *Proc. Int. Conf. Intell. Syst. Mol. Biol.* **1999**, 77–86 (1999).
27. Mintz, M., Bills, S., Snow, R. & Jurafsky, D. Distant supervision for relation extraction without labeled data. in *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, (eds Keh-Yih Su, Jian Su, Janyce Wiebe & Haizhou Li) 1003–1011 (Association for Computational Linguistics, 2009).
28. Blum, A. & Mitchell, T. M. Combining labeled and unlabeled data with co-training. in *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT 1998, Madison, Wisconsin, USA, July 24–26, 1998* (eds Bartlett, P. L. & Mansour, Y.) 92–100 (ACM, 1998).
29. Ma, Y., Cambria, E. & Gao, S. Label embedding for zero-shot fine-grained named entity typing. in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (eds Yuji Matsumoto & Rashmi Prasad). 171–180 (The COLING 2016 Organizing Committee, 2016).
30. Collins, M. & Singer, Y. Unsupervised models for named entity classification. in *1999 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora* (eds Pascale Fung & Joe Zhou) (Association for Computational Linguistics, 1999).
31. Medlock, B. & Briscoe, T. Weakly supervised learning for hedge classification in scientific literature. in *Proceedings of the 45th annual meeting of the association of computational linguistics* (eds Annie Zaenen & Antal van den Bosch). 992–999 (Association for Computational Linguistics, 2007).
32. Mann, G. S. & McCallum, A. Generalized expectation criteria for semi-supervised learning with weakly labeled data. *J. Mach. Learn. Res.* **11**, 955–984 (2010).
33. Khetan, A., Lipton, Z. C. & Anandkumar, A. Learning from noisy singly-labeled data. in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30–May 3, 2018, Conference Track Proceedings* (OpenReview.net, 2018).
34. Ratner, A., De Sa, C., Wu, S., Selsam, D. & Ré, C. Data programming: creating large training sets, quickly. *Adv. Neural Inf. Process. Syst.* **29**, 3567–3575 (2016).
35. Harkema, H., Dowling, J. N., Thornblade, T. & Chapman, W. W. Context: an algorithm for determining negation, experiencer, and temporal status from clinical reports. *J. Biomed. Inform.* **42**, 839–851 (2009).
36. Fries, J. A. Brundlefly at semeval-2016 task 12: recurrent neural networks vs. joint inference for clinical temporal information extraction. in *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16–17, 2016* (eds Bethard, S. et al.) 1274–1279 (The Association for Computer Linguistics, 2016).
37. Callahan, A. et al. Estimating the efficacy of symptom-based screening for COVID-19. *npj Digital Med.* **3**, 95 (2020).
38. Saeed, M. et al. Multiparameter intelligent monitoring in intensive care II (MIMIC-II): a public-access intensive care unit database. *Crit. Care Med.* **39**, 952 (2011).
39. Hanauer, D. Project EMERSE: COVID-19 synonyms. (2020) http://project-emerse.org/synonyms_covid19.html.
40. Rubin, D. L., Shah, N. H. & Noy, N. F. Biomedical ontologies: a functional perspective. *Brief. Bioinform.* **9**, 75–90 (2008).
41. Fu, S. et al. Clinical concept extraction: a methodology review. *J. Biomed. Inform.* **109**, 103526 (2020).
42. Friedman, C., Kra, P. & Rzhetsky, A. Two biomedical sublanguages: a description based on the theories of Zellig harris. *J. Biomed. Inform.* **35**, 222–235 (2002).
43. Wagner, T. et al. Augmented curation of clinical notes from a massive EHR system reveals symptoms of impending COVID-19 diagnosis. *eLife* **9**, e58227 <https://doi.org/10.7554/eLife.58227> (2020).
44. Wang, J., et al. COVID-19 SignSym: a fast adaptation of a general clinical NLP tool to identify and normalize COVID-19 signs and symptoms to OMOP common data model. *J. Am. Med. Inform. Assoc.* ocab015, (2021) <https://doi.org/doi.org.stanford.idm.oclc.org/10.1093/jamia/ocab015>
45. National Center for Advancing Translational Sciences (NCATS). National COVID cohort collaborative (N3C). <https://ncats.nih.gov/n3c> (2020).
46. Bach, S. H., He, B. D., Ratner, A. & Ré, C. Learning the structure of generative models without labeled data. in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6–11 August 2017*, Vol. 70 of *Proceedings of Machine Learning Research* (eds Precup, D. & Teh, Y. W.) 273–282 (PMLR, 2017).
47. Varma, P., Sala, F., He, A., Ratner, A. & Ré, C. Learning dependency structures for weak supervision models. in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA*, Vol. 97 of *Proceedings of Machine Learning Research* (eds Chaudhuri, K. & Salakhutdinov, R.) 6418–6427 (PMLR, 2019).
48. Honnibal, M., Montani, I., Van Landeghem, S. & Boyd, A. spaCy: industrial-strength natural language processing in Python. <https://doi.org/10.5281/zenodo.1212303> (2020).
49. UMLS® Reference Manual [Internet]. Bethesda (MD): National Library of Medicine (US); 2009 Sep-. 6, SPECIALIST Lexicon and Lexical Tools. Available from: <https://www.ncbi.nlm.nih.gov/books/NBK9680/>
50. Schriml, L. M. et al. Disease ontology: a backbone for disease semantic integration. *Nucleic Acids Res.* **40**, D940–6 (2012).
51. Degtyarenko, K. et al. ChEBI: a database and ontology for chemical entities of biological interest. *Nucleic Acids Res.* **36**, D344–50 (2008).
52. Davis, A. P. et al. The comparative toxicogenomics database’s 10th year anniversary: update 2015. *Nucleic Acids Res.* **43**, D914–20 (2015).
53. Zhou, W., Torvik, V. I. & Smalheiser, N. R. ADAM: another database of abbreviations in MEDLINE. *Bioinformatics* **22**, 2813–2818 (2006).
54. Wu, Y. et al. A long journey to short abbreviations: developing an open-source framework for clinical abbreviation recognition and disambiguation (CARD). *J. Am. Med. Inform. Assoc.* **24**, e79–e86 (2017).
55. McCray, A. T. An upper-level ontology for the biomedical domain. *Comp. Funct. Genom.* **4**, 80–84 (2003).
56. McCray, A. T., Burgun, A. & Bodenreider, O. Aggregating umls semantic types for reducing conceptual complexity. *Stud. Health Technol. Inform.* **84**, 216 (2001).
57. Schwartz, A. S. & Hearst, M. A. A simple algorithm for identifying abbreviation definitions in biomedical text. in *Proceedings of the 8th Pacific Symposium on Biocomputing, PSB*, (eds Altman, R. et al.) 451–462 (Pacific Symposium on Biocomputing, 2003).
58. Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Vol. 1 (Long and Short Papers)* (eds Jill Burstein, Christy Doran & Thamar Solorio), 4171–4186 (Association for Computational Linguistics, 2019).
59. Wei, J. & Zou, K. EDA: easy data augmentation techniques for boosting performance on text classification tasks. in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (eds Kentaro Inui, Jing Jiang, Vincent Ng & Xiaojun Wan) 6382–6388 (Association for Computational Linguistics, 2019).
60. Zhang, Y. & Yang, Q. A survey on multi-task learning. Preprint at <https://arxiv.org/abs/1707.08114> (2017).
61. Alsentzer, E. et al. Publicly available clinical BERT embeddings. in *Proceedings of the 2nd Clinical Natural Language Processing Workshop* (eds Anna Rumshisky, Kirk Roberts, Steven Bethard & Tristan Naumann), 72–78 (Association for Computational Linguistics, 2019).
62. Tjong Kim, E. & Buchholz, S. Introduction to the CONLL-2000 shared task: chunking. in *Proceedings of the Fourth Conference on Computational Natural Language Learning and of the Second Learning Language in Logic Workshop (CONLL/LL 2000)*. Lissabon, Portugal, 13–14 september 2000, 127–132 (ACL, 2000).
63. Wolf, T. et al. Transformers: state-of-the-art natural language processing. in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (eds Qun Liu & David Schlangen), 38–45 (Association for Computational Linguistics, 2020).
64. Fries, J. A. et al. Ontology-driven weak supervision for clinical entity classification in electronic health records. *Zenodo* <https://doi.org/10.5281/zenodo.4497214> (2021).
65. Dai, X., Karimi, S. & Paris, C. Medication and adverse event extraction from noisy text. in *Proceedings of the Australasian Language Technology Association*

- Workshop 2017* (eds Jojo Sze-Meng Wong & Gholamreza Haffari), 79–87 (Australasian Language Technology Association, 2017).
66. Si, Y., Wang, J., Xu, H. & Roberts, K. Enhancing clinical concept extraction with contextual embeddings. *J. Am. Med. Inform. Assoc.* **26**, 1297–1304 (2019).
 67. Lin, C., Dligach, D., Miller, T. A., Bethard, S. & Savova, G. K. Multilayered temporal modeling for the clinical domain. *J. Am. Med. Inform. Assoc.* **23**, 387–395 (2016).
 68. Mowery, D. L. et al. Task 2: ShARe/CLEF eHealth evaluation lab 2014. In *Working Notes for CLEF 2014 Conference, Sheffield, UK, September 15–18, 2014*, Vol. 1180 of *CEUR Workshop Proceedings* (eds Cappellato, L. et al.) 31–42 (CEUR-WS.org, 2014).
 69. Wei, C. et al. Assessing the state of the art in biomedical relation extraction: overview of the biocreative V chemical-disease relation (CDR) task. *Database J. Biol. Databases Curation* <https://doi.org/10.1093/database/baw032> (2016).
 70. Uzuner, O., Solti, I. & Cadag, E. Extracting medication information from clinical text. *J. Am. Med. Inform. Assoc.* **17**, 514–518 (2010).
 71. Bethard, S. et al. Semeval-2016 task 12: clinical tempeval. in *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016* (eds Bethard, S. et al.) 565–572 (The Association for Computer Linguistics, 2016).

Acknowledgements

This work was funded under NLM R01-LM011369-05. Thanks to Birju Patel and Keith Morse who did our COVID-19 clinical annotations and to Daisy Ding and Adrien Coulet who helped refine experimental hypotheses during the early stages of this project. Computational resources were provided by Nero, a shared big data computing platform made possible by the Stanford School of Medicine Research Office and Stanford Research Computing Center. Additional thanks to reader feedback from Stephen Pfohl, Erin Craig, Conor Corbin, and Jennifer Wilson.

Author contributions

J.A.F. conceived the initial study. J.A.F., E.S., S.K., S.F., J.P. and A.C. wrote code and conducted an experimental analysis of machine-learning models. A.C. and J.A.F. managed and adjudicated clinical text annotations. J.A.F., E.S. and N.H.S. contributed ideas

and experimental designs. N.H.S. supervised the project. All authors contributed to writing.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-021-22328-4>.

Correspondence and requests for materials should be addressed to J.A.F.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permission information is available at <http://www.nature.com/reprints>

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021