

Document-Level N -ary Relation Extraction with Multiscale Representation Learning

Robin Jia^{1*}

Cliff Wong²

Hoifung Poon²

¹ Stanford University, Stanford, California, USA

² Microsoft Research, Redmond, Washington, USA

robinjia@cs.stanford.edu, {cliff.wong, hoifung}@microsoft.com

Abstract

Most information extraction methods focus on binary relations expressed within single sentences. In high-value domains, however, n -ary relations are of great demand (e.g., drug-gene-mutation interactions in precision oncology). Such relations often involve entity mentions that are far apart in the document, yet existing work on cross-sentence relation extraction is generally confined to small text spans (e.g., three consecutive sentences), which severely limits recall. In this paper, we propose a novel multiscale neural architecture for document-level n -ary relation extraction. Our system combines representations learned over various text spans throughout the document and across the subrelation hierarchy. Widening the system’s purview to the entire document maximizes potential recall. Moreover, by integrating weak signals across the document, multiscale modeling increases precision, even in the presence of noisy labels from distant supervision. Experiments on biomedical machine reading show that our approach substantially outperforms previous n -ary relation extraction methods.

1 Introduction

Knowledge acquisition is a perennial challenge in AI. In high-value domains, it has acquired new urgency in recent years due to the advent of big data. For example, the dramatic drop in genome sequencing cost has created unprecedented opportunities for tailoring cancer treatment to a tumor’s genetic composition (Bahcall, 2015). Despite this potential, operationalizing personalized medicine is difficult, in part because it requires painstaking curation of precision oncology knowledge from biomedical literature. With tens of millions of papers on PubMed, and thousands more added every

“We next expressed *ALK F1174L*, *ALK F1174L/L1198P*, *ALK F1174L/G1123S*, and *ALK F1174L/G1123D* in the original *SH-SY5Y* cell line.”

(... 15 sentences spanning 3 paragraphs ...)

“The 2 mutations that were only found in the neuroblastoma resistance screen (*G1123S/D*) are located in the glycine-rich loop, which is known to be crucial for ATP and ligand binding and are the first mutations described that induce resistance to *TAE684*, but not to *PF02341066*.”

Figure 1: Two examples of drug-gene-mutation relations from a biomedical journal paper. The relations are expressed across multiple paragraphs, requiring document-level extraction.

day,¹ we are sorely in need of automated methods to accelerate manual curation.

Prior work in machine reading has made great strides in sentence-level binary relation extraction. However, generalizing extraction to n -ary relations poses new challenges. Higher-order relations often involve entity mentions that are far away in the document. Recent work on n -ary relation extraction has begun to explore cross-sentence extraction (Peng et al., 2017; Wang and Poon, 2018), but the scope is still confined to short text spans (e.g., three consecutive sentences), even though a document may contain hundreds of sentences and tens of thousands of words. While this already increases the yield compared to sentence-level extraction, it still misses many relations. For example, in Figure 1, the drug-gene-mutation relations between *PF02341066*, *ALK*, *G1123S(D)* (*PF02341066* can treat cancers with mutation *G1123S(D)* in gene *ALK*) can only be extracted by substantially expanding the scope. High-value information, such as latest medical findings, might only be mentioned once in the corpus. Maximiz-

*Work done as an intern at Microsoft Research.

¹ncbi.nlm.nih.gov/pubmed

ing recall is thus of paramount importance.

In this paper, we propose a novel multiscale neural architecture for document-level n -ary relation extraction. By expanding extraction scope to the entire document, rather than restricting relation candidates to co-occurring entities in a short text span, we ensure maximum potential recall. To combat the ensuing difficulties in document-level extraction, such as low precision, we introduce multiscale learning, which combines representations learned over text spans of varying scales and for various subrelations (Figure 2). This approach deviates from past methods in several key regards.

First, we adopt an entity-centric formulation by making a single prediction for each entity tuple occurring in a document. Previous n -ary relation extraction methods typically classify individual mention tuples, but this approach scales poorly to whole documents. Since each entity can be mentioned many times in the same document, applying mention-level methods leads to a combinatorial explosion of mention tuples. This creates not only computational challenges but also learning challenges, as the vast majority of these tuples do not express the relation. Our entity-centric formulation alleviates both of these problems.

Second, for each candidate tuple, prior methods typically take as input the contiguous text span encompassing the mentions. For document-level extraction, the resulting text span could become untenably large, even though most of it is unrelated to the relation of interest. Instead, we allow discontinuous input formed by multiple discourse units (e.g., sentence or paragraph) containing the given entity mentions.

Finally, while an n -ary relation might not reside within a discourse unit, its subrelations might. In Figure 1, the paper first mentions a gene-mutation subrelation, then discusses a drug-mutation subrelation in a later paragraph. By including subrelations in our modeling, we can predict n -ary relations even when all n entities never co-occur in the same discourse unit.

With multiscale learning, we turn the document view from a challenge into an advantage by combining weak signals across text spans and subrelations. Following recent work in cross-sentence relation extraction, we conduct thorough evaluation in biomedical machine reading. Our approach substantially outperforms prior n -ary relation extraction methods, attaining state-of-the-art results

on a large benchmark dataset recently released by a major cancer center. Ablation studies show that multiscale modeling is the key to these gains.²

2 Document-Level N -ary Relation Extraction

Prior work on relation extraction typically formulates it as a mention-level classification problem. Let e_1, \dots, e_n be entity mentions that co-occur in a text span T . Relation extraction amounts to classifying whether a relation R holds for e_1, \dots, e_n in T . For the well-studied case of binary relations within single sentences, $n = 2$ and T is a sentence.

In high-value domains, however, there is increasing demand for document-level n -ary relation extraction, where $n > 2$ and T is a full document that may contain hundreds of sentences. For example, a molecular tumor board needs to know if a drug is relevant for treating cancer patients with a certain mutation in a given gene. We can help the tumor board by extracting such ternary interactions from biomedical articles. The mention-centric view of relation extraction does not scale well to this general setting. Each of the n entities may be mentioned many times in a document, resulting in a large number of candidate mention tuples, even though the vast majority of them are irrelevant to the extraction task.

In this paper, we adopt an *entity-centric* formulation for document-level n -ary relation extraction. We use upper case for entities (E_1, \dots, E_n) and lower case for mentions (e_1, \dots, e_n). We define an n -ary relation candidate to be an $(n + 1)$ -tuple (E_1, \dots, E_n, T) , where each entity E_i is mentioned at least once in the text span T . The relation extraction model is given a candidate (E_1, \dots, E_n, T) and outputs whether or not the tuple expresses the relation R .³ Deciding what information to use from the various entity mentions within T is now a modeling question, which we address in the next section.

3 Our Approach: Multiscale Representation Learning

We present a general framework for document-level n -ary relation extraction using multiscale

² Our code and data will be available at hanover.azurewebsites.net

³ It is easy to extend our approach to situations where k mutually exclusive relations R_1, \dots, R_k must be distinguished, resulting in a $(k + 1)$ -way classification problem.

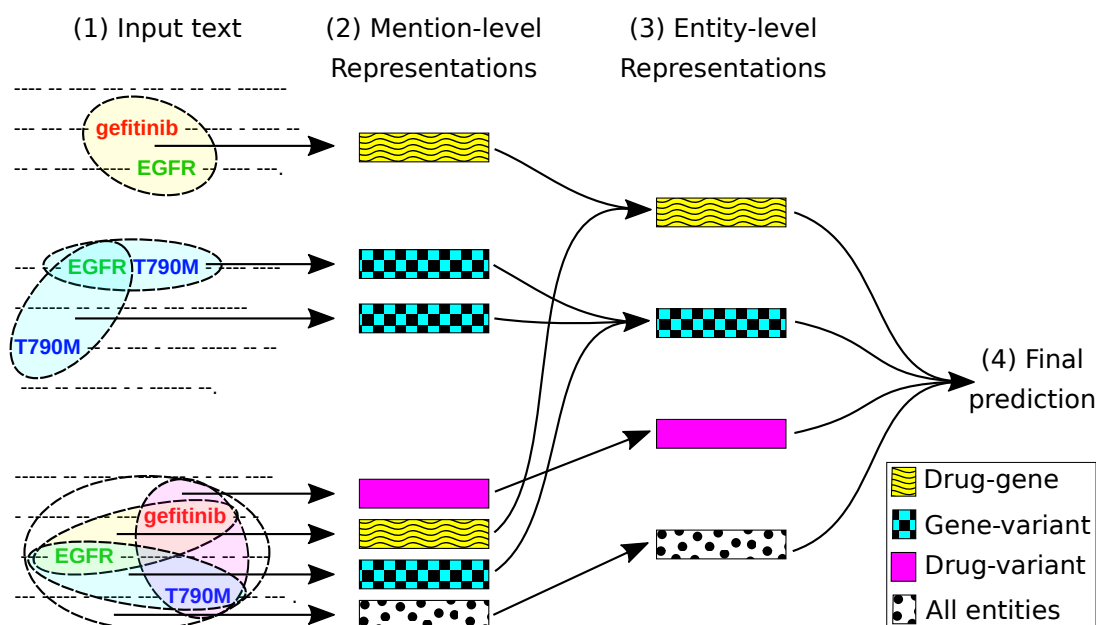


Figure 2: Multiscale representation learning for document-level n -ary relation extraction, an entity-centric approach that combines mention-level representations learned across text spans and subrelation hierarchy. (1) Entity mentions (e.g., **gefitinib**, a drug; **EGFR**, a gene; **T790M**, a variant) are identified from text, and mentions that co-occur within a discourse unit (e.g., paragraph) are isolated. (2) Within each discourse unit, mention-level representations are computed for each tuple of entity mentions. These representations may correspond to the entire n -ary relation or subrelations over subsets of entities (drug-variant, drug-gene, gene-variant). (3) At the document scale, mention-level representations for both the n -ary relation and its subrelations are combined into entity-level representations. (4) Entity-level representations are used to predict the relation.

representation learning. Given a document with text T and entities E_1, \dots, E_n , we first build mention-level representations for groups of these entities whenever they co-occur within the same discourse unit. We then aggregate these representations across the whole document, yielding entity-level representations for each subset of entities. Finally, we predict whether E_1, \dots, E_n participate in the relation based on the concatenation of these entity-level representations. These steps are depicted in Figure 2.

3.1 Mention-level Representation

Let the full document T be composed of discourse units T_1, \dots, T_m (e.g., different paragraphs). Let T_j be one such discourse unit, and suppose e_1, \dots, e_n are entity mentions of E_1, \dots, E_n that co-occur in T_j . We construct a contextualized representation for mention tuple (e_1, \dots, e_n) in T_j . In this paper, we use a standard approach by applying a bi-directional LSTM (BiLSTM) to T_j , concatenating the hidden states for each mention, and feeding this through a single-layer neural network. We denote the resulting vector as $r(R, e_1, \dots, e_n, T_j)$ for the relation R .

3.2 Entity-level Representation

Let $M(R, E_1, \dots, E_n, T)$ denote the set of all mention tuples (e_1, \dots, e_n) and discourse units T_j within T such that each e_i appears in T_j . We can create an entity-level representation $r(R, E_1, \dots, E_n, T)$ of the n entities by combining mention-level representations using an aggregation operator \mathcal{C} :

$$r(R, E_1, \dots, E_n, T) = \mathcal{C}_{(e_1, \dots, e_n, T_j) \in M(R, E_1, \dots, E_n, T)} r(R, e_1, \dots, e_n, T_j)$$

A standard choice for \mathcal{C} is max pooling, which works well if it is pretty clear-cut whether a mention tuple expresses a relation. In practice, however, the mention tuples could be ambiguous and less than certain individually, yet collectively express a relation in the document. This motivates us to experiment with logsumexp, the smooth version of max, where

$$\text{logsumexp}(x_1, \dots, x_k) = \log \sum_{i=1}^k \exp(x_i).$$

This facilitates accumulating weak signals from individual mention tuples, and our experiments show that it substantially improves extraction accuracy compared to max pooling.

3.3 Subrelations

For higher-order relations (i.e., larger n), it is less likely that they will be completely contained within a discourse unit. Often, the relation can be decomposed into subrelations over subsets of entities, each of which is more likely to be expressed in a single discourse unit. This motivates us to construct entity-level representations for subrelations as well. The process is straightforward. Let R_S be the $|S|$ -ary subrelation over entities $E_{S_1}, \dots, E_{S_{|S|}}$, where $S \subseteq \{1, \dots, n\}$ and $|S|$ denotes its size. We first construct mention-level representations $r(R_S, e_{S_1}, \dots, e_{S_{|S|}}, T)$ for R_S and its relevant entity mentions, then combine them into an entity-level representation $r(R_S, E_{S_1}, \dots, E_{S_{|S|}}, D)$ using the chosen aggregation operator \mathcal{C} . We do this for every $S \subseteq \{1, \dots, n\}$ with $|S| \geq 2$ (including the whole set, which corresponds to the full relation R). This gives us an entity-level representation for each subrelation of arity at least 2, or equivalently, each subset of entities of size at least 2.

3.4 Relation Prediction

To make a final prediction, we first concatenate all of the entity-level representations $r(R_S, E_{S_1}, \dots, E_{S_{|S|}}, D)$ for all $S \subseteq \{1, \dots, n\}$ with $|S| \geq 2$. The concatenated representation is fed through a two-layer feedforward neural network followed by a softmax function to predict the relation type.

It is possible that for some subrelations R_S , all $|S|$ entities do not co-occur in any discourse unit. When this happens, we set $r(R_S, E_{S_1}, \dots, E_{S_{|S|}})$ to a bias vector which is learned separately for each R_S . This ensures that the concatenation is done over a fixed number of vectors, e.g., 4 for a ternary relation (three binary subrelations and the main relation). Importantly, this strategy enables us to make meaningful predictions for relation candidates even if all n entities never co-occur in the same discourse unit; such candidates would never be generated by a system that only looks at single discourse units in isolation.

3.5 Document Model

Our document model is actually a family of representation learning methods, conditioned on the choice of discourse units, subrelations, and aggregation operators. In this paper, we consider sentences and paragraphs as possible discourse

	Sentence level	Paragraph level	Document level
Text Units	2,326	3,687	3,362
Pos. Examples	2,222	4,906	8,514
Neg. Examples	2,849	13,371	323,584

Table 1: Statistics of our training corpus using PMC-OA articles and distant supervision from CIVIC, GDKD, and OncoKB. “Text Units” refers to the number of distinct sentences, paragraphs, and documents that contain a candidate triple of drug, gene, mutation.

	Development	Test
Documents	118	225
Annotated facts	701	1,324
Paragraphs per document	101	105
Sentences per document	314	320
Words per document	6,871	7,010

Table 2: Statistics of the CKB evaluation corpus.

units. We explore max and logsumexp as aggregation operators. Moreover, we explore ensemble prediction as an additional aggregation method. Specifically, we learn a restricted multiscale model by limiting the text span to a single discourse unit (e.g., a paragraph); the model still combines representations across mentions and subrelations. At test time, given a full document with m discourse units, we obtain independent predictions p_1, \dots, p_m for each discourse unit. We then combine these probabilities using an ensemble operator \mathcal{P} . A natural choice for \mathcal{P} is max, though we also experiment with noisy-or:

$$\mathcal{P}(p_1, \dots, p_k) = 1 - \prod_{i=1}^k (1 - p_i).$$

It is also possible to ensemble multiple models that operate on different discourse units, using this same operator.

Our model can be trained using standard supervised or indirectly supervised methods. In this paper, we focus on distant supervision, as it is a particularly potent learning paradigm for high-value domains. Our entity-centric formulation is particularly well aligned with distant supervision, as distant supervision at the entity level is significantly less noisy compared to the mention level, so we don’t need to deploy sophisticated denoising strategies such as multi-instance learning (Hoffmann et al., 2011).

4 Experiments

4.1 Biomedical Machine Reading

We validate our approach on a standard biomedical machine reading task: extracting drug-gene-mutation interactions from biomedical articles (Peng et al., 2017; Wang and Poon, 2018). We cast this task as binary classification: given a drug, gene, mutation, and document in which they are mentioned, determine whether the document asserts that the mutation in the gene affects response to the drug. For training, we use documents from the PubMed Central Open Access Subset (PMCOA)⁴. For distant supervision, we use three existing knowledgebases (KBs) with hand-curated drug-gene-mutation facts: CIVIC,⁵ GDKD (Dienstmann et al., 2015), and OncoKB (Chakravarty et al., 2017). Table 1 shows basic statistics of this training data. Past methods using distant supervision often need to up-weight positive examples, due to the large proportion of negative candidates. Interestingly, we found that our document model was robust to this imbalance, as re-weighting had little effect and we didn't use it in our final results.

Evaluating distant supervision methods is challenging, as there is often no gold-standard test set, especially at the mention level. Prior work thus resorts to reporting sample precision (estimated proportion of correct system extractions) and absolute recall (estimated number of correct system extractions). This requires subsampling extraction results and manually annotating them. Subsampling variance also introduces noise in the estimate.

Instead, we used CKB CORE™, a public subset of the Clinical Knowledgebase (CKB)⁶ (Patterson et al., 2016), as our gold-standard test set. CKB CORE™ contains document-level annotation of drug-gene-mutation interactions manually curated by The Jackson Laboratory (JAX), an NCI-designated cancer center. It is a high-quality KB containing facts from a few hundred PubMed articles for 86 genes, with minimal overlap with the three KBs we used for distant supervision. To avoid contamination, we removed CKB entries whose documents were used in our training data, and split the rest into a development and test set. See Table 2 for statistics. We tuned hyperparameters and thresholds on the development set, and report results on the test set.

⁴www.ncbi.nlm.nih.gov/pmc

⁵civic.genome.wustl.edu

⁶ckbhome.jax.org

4.2 Implementation Details

We conducted standard preprocessing and entity linking, similar to Wang and Poon (2018) (see Section A.1). Following standard practice, we masked all entities of the same type with a dummy token, to prevent the classifier from simply memorizing the facts in distant supervision. Wang and Poon (2018) observed that many errors stemmed from incorrect gene-mutation association. We therefore developed a simple rule-based system that predicts which gene-mutation pairs are valid (see Section A.2). We removed candidates that contained a gene-mutation pair that was not predicted by the rule-based system.

4.3 Main Results

We evaluate primarily on area under the precision recall curve (AUC).⁷ We also report maximum recall, which is the fraction of true facts for which a candidate was generated. Finally, we report precision, recall, and F1, using a threshold tuned to maximize F1 on the CKB development set.

We compared our multiscale system (MULTISCALE) with three restricted variants (SENTLEVEL, PARALEVEL, DOCLEVEL). SENTLEVEL and PARALEVEL restricted training and prediction to single discourse units (i.e., sentences and paragraphs), and produced a document-level prediction by applying the ensemble operator over individual discourse units. DOCLEVEL takes the whole document as input, with each paragraph as a discourse unit. MULTISCALE further combined SENTLEVEL, PARALEVEL, and DOCLEVEL using the ensemble operator. For additional details about the models, see Section A.3. We also compared MULTISCALE with DPL (Wang and Poon, 2018), the prior state of the art in cross-sentence n -ary relation extraction. DPL classifies drug-gene-mutation interactions within three consecutive sentences using the same model architecture as Peng et al. (2017), but incorporates additional indirect supervision such as data programming and joint inference. We used the DPL code from the authors and produced a document-level prediction similarly using the ensemble operator. In the base version, we used max as the ensemble operator. We also evaluated the effect when we

⁷ We compute area using average precision, which is similar to a right Riemann sum. This avoids errors introduced by the trapezoidal rule, which may overestimate area.

System	AUC	Max Recall	Precision	Recall	F1
Base versions					
DPL	24.4	53.8	27.3	42.3	33.2
SENTLEVEL	22.4	36.6	39.3	34.7	36.9
PARALEVEL	33.1	58.9	36.5	44.6	40.1
DOCLEVEL	36.7	79.0	45.4	38.5	41.7
MULTISCALE	37.3	79.0	41.8	43.3	42.5
+ Noisy-Or					
DPL	31.5	53.8	33.3	41.5	36.9
SENTLEVEL	25.3	36.6	39.3	35.3	37.2
PARALEVEL	35.6	58.9	44.3	40.6	42.4
DOCLEVEL	36.7	79.0	45.4	38.5	41.7
MULTISCALE	39.7	79.0	48.1	38.9	43.0
+ Noisy-Or + Gene-mutation filter					
DPL	39.1	52.6	50.5	47.8	49.1
SENTLEVEL	29.0	35.5	63.3	34.2	44.4
PARALEVEL	42.1	57.2	50.6	50.7	50.7
DOCLEVEL	42.9	74.4	49.3	46.6	47.9
MULTISCALE	47.5	74.4	52.6	53.0	52.8

Table 3: Comparison of our multiscale system with restricted variants and DPL (Wang and Poon, 2018) on CKB.

used noisy-or as the ensemble operator, as well as when we applied the gene-mutation filter during postprocessing.

Table 3 shows the results on the CKB test set. In all scenarios, our full model (MULTISCALE) substantially outperforms the prior state-of-the-art system (DPL). For example, in the best setting, using both noisy-or and the gene-mutation filter, the full model improves over DPL by 8.4 AUC points. Multiscale learning is the key to this performance gain, with MULTISCALE substantially outperforming more restricted variants. Not surprisingly, expanding extraction scope from sentences to paragraphs resulted in the biggest gain, already surpassing DPL. Conducting end-to-end learning over a document-level representation, as in DOCLEVEL, is beneficial compared to ensembling over predictions for individual discourse units (SENTLEVEL, PARALEVEL), especially in the base version. Interestingly, MULTISCALE still attained significant gain over DOCLEVEL with an ensemble over SENTLEVEL and PARALEVEL, suggesting that the document-level representation can still be improved. In addition to prediction accuracy, the document-level models also have much more room to grow, as maximum recall is about 20 absolute points higher in MULTISCALE and DOCLEVEL, compared to PARALEVEL or DPL.⁸

The ensemble operator had a surprisingly large effect, as shown by the gain when it was changed from max (base version) to noisy-or. This sug-

⁸The difference in actual recall is less pronounced, as we chose thresholds to maximize F1 score. We expect actual recall to increase significantly as document-level models improve, whereas the other models are closer to their ceiling.




Figure 3: Precision-recall curves on CKB (with noisy-or and gene-mutation filter). MULTISCALE attained generally better precision than PARALEVEL, and higher maximum recall like DOCLEVEL.

gests that combining weak signals across multiple scales can be quite beneficial. Our handcrafted gene-mutation filter also improved all systems substantially, corroborating the analysis of Wang and Poon (2018). In particular, without the filter, it is hard for the document-level models to achieve high precision, so they sacrifice a lot of recall to get good F1 scores. Using the filter helps them attain significantly higher recall while maintaining respectable precision.

Figure 3 shows the precision-recall curves for the four models (with noisy-or and gene-mutation filter). DOCLEVEL has higher maximum recall than PARALEVEL, but generally lower precision at the same recall level. By ensembling all three variants, MULTISCALE achieves the best combi-

System	AUC	MR	P	R	F1
MULTISCALE	47.5	74.4	52.6	53.0	52.8
- SENTLEVEL	47.0	74.4	43.0	55.8	48.6
- PARALEVEL	45.9	74.4	48.8	49.8	49.3
- DOCLEVEL	42.4	57.2	59.6	44.4	50.9

Table 4: Results on CKB when removing either SENTLEVEL, PARALEVEL, or DOCLEVEL from the ensemble computed by MULTISCALE. MR=max recall, P=precision, R=recall.

System	AUC	P	R	F1
SENTLEVEL	28.3	62.7	35.1	45.0
PARALEVEL	38.1	47.4	52.2	49.7
DOCLEVEL	41.1	48.2	45.6	46.9
MULTISCALE	43.7	45.7	51.2	48.3

Table 5: Results on CKB after replacing logsumexp with max (with noisy-or and gene-mutation filter). P=precision, R=recall. Max recall same as in Table 3.

nation: it generally improves precision while capturing more cross-paragraph relations. This can also be seen in Table 4, where we ablate each of the three variants used by MULTISCALE. All three variants in the ensemble contributed to overall performance.

We use logsumexp as the aggregation operator to combine mention-level representations into an entity-level one. If we replace it with max pooling, the performance drops substantially across the board, as shown in Table 5. For example, MULTISCALE lost 3.8 absolute points in AUC. Such difference is also observed in Verga et al. (2018). As in comparing ensemble operators, this demonstrates the benefit of combining weak signals using a multiscale representation.

4.4 Cross-sentence and Cross-paragraph Extractions

Compared to standard sentence-level extraction, our method can extract relations among entities that never co-occur in the same sentence or even paragraph. Figure 4 shows the proportion of correctly predicted facts by MULTISCALE that are expressed across paragraph or sentence boundaries. MULTISCALE can substantially improve the recall by making additional cross-sentence and cross-paragraph extractions. We manually inspected twenty correct cross-paragraph extractions (with the chosen threshold for the precision/recall numbers in Table 3) and found that our model was able to handle some interesting linguistic phenomena. Often, a paper would first describe the mutations present in a patient cohort,

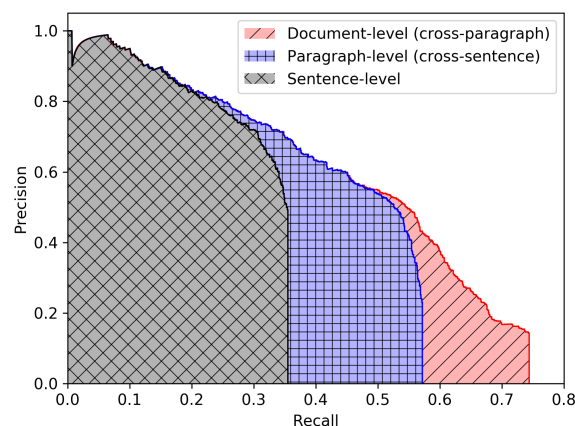


Figure 4: Breakdown of MULTISCALE recall based on whether entities in a correctly extracted fact occurred within a single sentence, cross-sentence but within a single paragraph, or only cross-paragraph. Adding cross-sentence and cross-paragraph extractions is important for high recall.

System	AUC	MR	P	R	F1
Base version					
SENTDRUGMUT	31.0	40.8	60.0	40.7	48.5
SENTDRUGGENE	17.9	64.2	31.4	27.9	29.5
PARADRUGMUT	39.9	57.7	49.3	50.3	49.8
PARADRUGGENE	19.9	68.9	32.1	18.9	23.8
+ Noisy-Or					
SENTDRUGMUT	32.6	40.8	61.1	39.4	47.9
SENTDRUGGENE	23.5	64.2	36.3	34.2	35.2
PARADRUGMUT	42.0	57.7	49.9	51.5	50.7
PARADRUGGENE	26.1	68.9	46.1	29.5	36.0

Table 6: Results of subrelation decomposition baselines on CKB, with the gene-mutation filter. MR=max recall, P=precision, R=recall.

and later describe the effects of drug treatment. There are also instances of bridging anaphora, for example via cell lines. One paper first stated the gene and mutation for a cell line “*The FLT3-inhibitor resistant cells Ba/F3-ITD+691, Ba/F3-ITD+842, ..., which harbored FLT-ITD plus F691L, Y842C, ... mutations...*”, and later stated the drug effect on the cell line “*E6201 also demonstrated strong anti-proliferative effects in FLT3-inhibitor resistant cells... such as Ba/F3-ITD+691, Ba/F3-ITD+842...*”.

4.5 Subrelation Decomposition

As a baseline, we also consider a different document-level strategy where we decompose the n -ary relation into subrelations of lower arity, train independent classifiers for them, then join the subrelation predictions into one for the n -ary relation. We found that with distant supervision, the

gene-mutation subrelation classifier was too noisy. Therefore, we focused on training drug-gene and drug-mutation classifiers, and joined each with the rule-based gene-mutation predictions to make ternary predictions. Table 6 shows the results on CKB. The paragraph-level drug-mutation model is quite competitive, which benefits from the fact that the gene-mutation associations in a document are unique. This is not true in general n -ary relations. Still, it trails MULTISCALE by a large margin in predictive accuracy, and with an even larger gap in the potential upside (i.e., maximum recall). The drug-gene model has higher maximum recall, but much worse precision. This low precision is expected, as it is usually not valid to assume that if a drug and gene interact, then all possible mutations in the gene will have an effect on the drug response.

4.6 Error Analysis

While much higher compared to other systems, the maximum recall for MULTISCALE is still far from 100%. For over 20% of the relations, we can't find all three entities in the document. In many cases, the missing entities are in figures or supplements, beyond the scope of our extraction. Some mutations are indirectly referenced by well-known cell lines. There are also remaining entity linking errors (e.g., due to missing drug synonyms).

We next manually analyzed some sample prediction errors. Among 50 false positive errors, we found a significant portion of them were actually true mentions in the paper but were excluded by curators due to additional curation criteria. For example, CKB does not curate a fact referenced in related work, or if they deem the empirical evidence as insufficient. This suggests the need for even higher-order relation extraction to cover these aspects. We also inspected 50 sample false negative errors. In 40% of the cases, the textual evidence is vague and requires corroboration from a table or figure. In most of the remaining cases, there is direct textual evidence, though they require cross-paragraph reasoning (e.g., bridging anaphora). While MULTISCALE was able to process such phenomena sometimes, there is clearly much room to improve.

5 Related Work

N -ary relation extraction Prior work on n -ary relation extraction generally follows Davidsonian

semantics by reducing the n -ary relation to n binary relations between the reified relation and its arguments, a.k.a. *slot filling*. For example, early work on the Message Understanding Conference (MUC) dataset aims to identify event participants in news articles (Chinchor, 1998). More recently, there has been much work in extracting semantic roles for verbs, as in semantic role labeling (Palmer et al., 2010), as well as properties for popular entities, as in Wikipedia Infobox (Wu and Weld, 2007) and TAC KBP⁹. In biomedicine, the BioNLP Event Extraction Shared Task aims to extract genetic events such as expression and regulation (Kim et al., 2009). These approaches typically assume that the whole document refers to a single coherent event, or require an event anchor (e.g., verb in semantic role labeling and trigger word in event extraction). We instead follow recent work in cross-sentence n -ary relation extraction (Peng et al., 2017; Wang and Poon, 2018; Song et al., 2018), which does not have these restrictions.

Document-level relation extraction Most information extraction work focuses on modeling and prediction within sentences (Surdeanu and Ji, 2014). Duan et al. (2017) introduces a pre-trained document embedding to aid event detection, but their extraction is still at the sentence level. Past work on cross-sentence extraction often relies on explicit coreference annotations or the assumption of a single event in the document (Wick et al., 2006; Gerber and Chai, 2010; Swampilalai and Stevenson, 2011; Yoshikawa et al., 2011; Koch et al., 2014; Yang and Mitchell, 2016). Recently, there has been increasing interest in general cross-sentence relation extraction (Quirk and Poon, 2017; Peng et al., 2017; Wang and Poon, 2018), but their scope is still limited to short text spans of a few consecutive sentences. These methods all extract relations at the mention level, which does not scale to whole documents due to the combinatorial explosion of relation candidates. Wu et al. (2018b) applies manually crafted rules to heavily filter the candidates. We instead adopt an entity-centric approach and combine mention-level representations to create an entity-level representation for extraction. Mintz et al. (2009) aggregates mention-level features into entity-level ones within a document, but they only consider

⁹<http://www.nist.gov/tac/2016/KBP/ColdStart/index.html>

binary relations within single sentences. Kilicoglu (2016) used hand-crafted features to improve cross-sentence extraction, but they focus on binary relations, and their documents are limited to abstracts, which are substantially shorter than the full-text articles we consider. Verga et al. (2018) applies self-attention to combine the representations of all mention pairs into an entity pair representation, which can be viewed a special case of our framework. Their work is also limited to binary relations and abstracts, rather than full documents.

Multiscale modeling Deep learning on long sequences can benefit from multiscale modeling that accounts for varying scales in the discourse structure. Prior work focuses on generative learning such as language modeling (Chung et al., 2017). We instead apply multiscale modeling to discriminative learning for relation extraction. In addition to modeling various scales of discourse units (sentence, paragraph, document), we also combine mention-level representations into an entity-level one, as well as sub-relations of the n -ary relation. McDonald et al. (2005) learn $\binom{n}{2}$ pairwise relation classifiers, then construct maximal cliques of related entities, which also bears resemblance to our subrelation modeling. However, our approach incorporates the entire subrelation hierarchy, provides a principled end-to-end learning framework, and extracts relations from the whole document rather than within single sentences.

Distant supervision Distant supervision has emerged as a powerful paradigm to generate large but potentially noisy labeled datasets (Craven et al., 1999; Mintz et al., 2009). A common denoising strategy applies multi-instance learning by treating mention-level labels as latent variables (Hoffmann et al., 2011). Noise from distant supervision increases as extraction scope expands beyond single sentences, motivating a variety of indirect supervision approaches (Quirk and Poon, 2017; Peng et al., 2017; Wang and Poon, 2018). Our entity-centric representation and multiscale modeling provide an orthogonal approach to combat noise by combining weak signals spanning various text spans and subrelations.

6 Conclusion

We propose a multiscale, entity-centric approach for document-level n -ary relation extraction.

We vastly increase maximum recall by scoring document-level candidates. Meanwhile, we preserve precision with a multiscale approach that combines representations learned across the subrelation hierarchy and text spans of various scales. Our method substantially outperforms prior cross-sentence n -ary relation extraction approaches in the high-value domain of precision oncology.

Our document-level view opens opportunities for multimodal learning by integrating information from tables and figures (Wu et al., 2018a). We used the ternary drug-gene-mutation relation as a running example in this paper, but knowledge bases often store additional fields such as effect (sensitive or resistance), cancer type (solid tumor or leukemia), and evidence (human trial or cell line experiment). It is straightforward to apply our method to such higher-order relations. Finally, it will be interesting to validate our approach in a real-world assisted-curation setting, where a machine reading system proposes candidate facts to be verified by human curators.

Acknowledgements

We thank Sara Patterson and Susan Mockus for guidance on precision oncology knowledge curation and CKB data, Hai Wang for help in running experiments with deep probabilistic logic, and Tristan Naumann, Rajesh Rao, Peng Qi, John Hewitt, and the anonymous reviewers for their helpful comments. R.J. is supported in part by an NSF Graduate Research Fellowship under Grant No. DGE-114747.

References

- Orli Bahcall. 2015. Precision medicine. *Nature*, 526:335.
- Debyani Chakravarty, Jianjiong Gao, Sarah Phillips, Ritika Kundra, Hongxin Zhang, Jiaojiao Wang, Julia E. Rudolph, Rona Yaeger, Tara Soumerai, Moriah H. Nissan, Matthew T. Chang, Sarat Chandrapaty, Tiffany A. Traina, Paul K. Paik, Alan L. Ho, Feras M. Hantash, Andrew Grupe, Shrujal S. Baxi, Margaret K. Callahan, Alexandra Snyder, Ping Chi, Daniel C. Danila, Mrinal Gounder, James J. Harding, Matthew D. Hellmann, Gopa Iyer, Yelena Y. Janjigian, Thomas Kaley, Douglas A. Levine, Maevae Lowery, Antonio Omuro, Michael A. Postow, Dana Rathkopf, Alexander N. Shoushtari, Neerav Shukla, Martin H. Voss, Ederlinda Paraiso, Ahmet Zehir, Michael F. Berger, Barry S. Taylor, Leonard B. Saltz, Gregory J. Riely, Marc Ladanyi, David M. Hyman, Jos Baselga, Paul Sabbatini, David B. Solit,

- and Nikolaus Schultz. 2017. Oncokb: A precision oncology knowledge base. *JCO Precision Oncology*.
- Nancy Chinchor. 1998. Overview of MUC-7/MET-2. Technical report, Science Applications International Corporation, San Diego, CA.
- Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. 2017. Hierarchical multiscale recurrent neural networks. In *ICLR*.
- M. Craven, J. Kumlien, et al. 1999. Constructing biological knowledge bases by extracting information from text sources. In *ISMB*, pages 77–86.
- Rodrigo Dienstmann, In Sock Jang, Brian Bot, Stephen Friend, and Justin Guinney. 2015. Database of genomic biomarkers for cancer drugs and clinical targetability in solid tumors. *Cancer Discovery*, 5.
- Shaoyang Duan, Ruifang He, and Wenli Zhao. 2017. Exploiting document level information to improve event detection via recurrent neural networks. In *IJCNLP*.
- Matthew Gerber and Joyce Y. Chai. 2010. Beyond NomBank: A study of implicit arguments for nominal predicates. In *Proceedings of the Forty-Eighth Annual Meeting of the Association for Computational Linguistics*.
- R. Hoffmann, C. Zhang, X. Ling, L. S. Zettlemoyer, and D. S. Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Association for Computational Linguistics (ACL)*, pages 541–550.
- Halil Kilicoglu. 2016. Inferring implicit causal relationships in biomedical literature. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*.
- Jin-Dong Kim, Tomoko Ohta, Sampo Pyysalo, Yoshinobu Kano, and Junichi Tsujii. 2009. Overview of bionlp-09 shared task on event extraction. In *Proceedings of the BioNLP Workshop 2009*.
- D. Kingma and J. Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Mitchell Koch, John Gilmer, Stephen Soderland, and S. Daniel Weld. 2014. Type-aware distantly supervised relation extraction with linked arguments. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1891–1901. Association for Computational Linguistics.
- Ryan McDonald, Fernando Pereira, Seth Kulick, Scott Winters, Yang Jin, and Pete White. 2005. Simple algorithms for complex relation extraction with applications to biomedical IE. In *Proceedings of the Forty-Third Annual Meeting on Association for Computational Linguistics*.
- M. Mintz, S. Bills, R. Snow, and D. Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Association for Computational Linguistics (ACL)*, pages 1003–1011.
- Martha Palmer, Daniel Gildea, and Nianwen Xue. 2010. Semantic role labeling. *Synthesis Lectures on Human Language Technologies*, 3(1).
- Sarah E. Patterson, Rangjiao Liu, Cara M. Statz, Daniel Durkin, Anuradha Lakshminarayana, and Susan M. Mockus. 2016. The clinical trial landscape in oncology and connectivity of somatic mutational profiles to targeted therapies. *Human Genomics*, 10.
- Nanyun Peng, Hoifung Poon, Chris Quirk, and Kristina Toutanova Wen tau Yih. 2017. Cross-sentence *N*-ary relation extraction with graph LSTMs. *Transactions of the Association for Computational Linguistics*, 5:101–115.
- Sampo Pyysalo, F Ginter, Hans Moen, T Salakoski, and Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing. In *Proceedings of Languages in Biology and Medicine*.
- Chris Quirk and Hoifung Poon. 2017. Distant supervision for relation extraction beyond the sentence boundary. In *Proceedings of the Fifteenth Conference on European chapter of the Association for Computational Linguistics*.
- Lin Feng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. *N*-ary relation extraction using graph state LSTM. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Mihai Surdeanu and Heng Ji. 2014. Overview of the English Slot Filling Track at the TAC2014 Knowledge Base Population evaluation. In *Proceedings of the TAC-KBP 2014 Workshop*, pages 1–15, Gaithersburg, Maryland, USA.
- Kumutha Swampillai and Mark Stevenson. 2011. Extracting relations within and across sentences. In *Proceedings of the Conference on Recent Advances in Natural Language Processing*.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Patrick Verga, Emma Strubell, and Andrew McCallum. 2018. Simultaneously self-attending to all mentions for full-abstract biological relation extraction. In *NAACL*.
- Hai Wang and Hoifung Poon. 2018. Deep probabilistic logic: A unifying framework for indirect supervision. In *EMNLP*.
- Michael Wick, Aron Culotta, and Andrew McCallum. 2006. Learning field compatibilities to extract

database records from unstructured text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Fei Wu and Daniel S. Weld. 2007. Autonomously semantifying wikipedia. In *Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management*.

S. Wu, L. Hsiao, X. Cheng, B. Hancock, T. Rekatsinas, P. Levis, and C. R'e. 2018a. Fonduer: Knowledge base construction from richly formatted data. In *Proceedings of SIGMOD 2018*.

Sen Wu, Luke Hsiao, Xiao Cheng, Braden Hancock, Theodoros Rekatsinas, Philip Levis, and Christopher R. 2018b. Fonduer: Knowledge base construction from richly formatted data. In *SIGMOD*.

Bishan Yang and Tom Mitchell. 2016. Joint extraction of events and entities within a document context. In *NAACL*.

Katsumasa Yoshikawa, Sebastian Riedel, Tsutomu Hirao, Masayuki Asahara, and Yuji Matsumoto. 2011. Coreference based event-argument relation extraction on biomedical text. *Journal of Biomedical Semantics*, 2(5).

A Appendices

A.1 Preprocessing

Full-text documents in this study were obtained from PMC. The text was first tokenized using NLTK¹⁰, then entities were extracted using a combination of regular expressions and dictionary lookups. To identify mutation mentions, we applied a regular expression rule for missense mutations. To identify gene mentions, we used dictionary lookup from the HUGO Gene Nomenclature Committee (HGNC)¹¹ dataset. To identify drug mentions, we used dictionary lookup from a curated list of drugs and their synonyms. For our training set, our list of drugs consists of all the drugs present in the distant supervision knowledge bases and selected cancer-related drugs from DrugBank¹² (770 drugs total). For our test set, our drug dictionary consists of all the drugs in CKB (1119 drugs).

A.2 Gene-mutation Rule-based System

Here we describe our rule-based system for linking mutations and genes within a document. We first generate a global mapping of mutations to sets of genes by combining publicly-available

mutation-gene datasets (COSMIC¹³, COSMIC Cell Lines Project¹⁴, CIViC¹⁵, and OncoKB¹⁶). We then augment this mapping by finding the gene that most frequently co-occurs with each mutation in all of PubMed Central (PMC) full-text articles based on three high-precision rules:

1. Gene and mutation are in the same token (e.g., "EGFR-T790M")
2. Gene token is followed by mutation token (e.g., "EGFR T790M")
3. Gene token is followed by a token of any single character and then followed by mutation token (e.g., "EGFR - T790M")

For each mutation, we start with the first rule, and find all text matches for a gene with that mutation and rule. If we found at least one match, we add the gene that occurred in the most matches to the global map. Otherwise, we repeat with the next rule.

Each mutation in the global mutation-gene map is mapped to more than 20 genes on average. However, within the context of a document, each mutation is (usually) associated with just a single gene. Given a document containing a mutation, we associate that mutation with the gene that (1) is in the global mutation-gene map for that mutation, and (2) appears closest to any mutation mention in the document.

To associate genes for the remaining mutations, we apply two recall-friendly regular expression rules within that document:

4. Mutation is in same sentence as "*GENE* mut"
5. Mutation is in same paragraph as "*GENE* mutation"

We choose the first gene in the document that satisfies one of the two rules, in the above order. If there is still no matching gene at this point, the most frequent gene in the document is selected for that mutation.

A.3 Model Details

We used 200-dimensional word vectors, initialized with word2vec vectors trained on a biomedical text corpus (Pyysalo et al., 2013). We updated

¹⁰<https://www.nltk.org/>

¹¹<https://www.genenames.org/>

¹²<https://www.drugbank.ca/>

¹³<https://cancer.sanger.ac.uk/cosmic>

¹⁴https://cancer.sanger.ac.uk/cell_lines

¹⁵<https://civcdb.org>

¹⁶<http://oncokb.org/>

these vectors during training. At each step, our BiLSTM received as input a concatenation of the word vectors and a 100-dimensional embedding of the index of the current discourse unit within the document. Following [Vaswani et al. \(2017\)](#), we used sinusoidal embeddings to represent these indices. We used a single-layer bidirectional LSTM with a 200-dimensional hidden state. Mention-level representations were 400-dimensional and computed from BiLSTM hidden states using a single linear layer followed by the tanh activation function. For the final prediction layer, we used a two-layer feedforward network with 400 hidden units and ReLU activation function. We train using the Adam optimizer ([Kingma and Ba, 2014](#)) with learning rate of 1×10^{-5} . During training, we consider each document to be a single batch, which allows us to reuse computation for different relation candidates in the same document.