

## Representation and Normalization of Complex Interventions for Evidence Computing

Zhehuan Chen<sup>a</sup>, Hao Liu<sup>a</sup>, Stan Liao<sup>b</sup>, Marguerite Bernard<sup>c</sup>, Tian Kang<sup>a</sup>, Latoya A Stewart<sup>d</sup>,  
Chunhua Weng<sup>a</sup>

<sup>a</sup>Department of Biomedical Informatics, Columbia University, New York, NY, USA

<sup>b</sup>Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY, USA

<sup>c</sup>Institute of Human Nutrition, Columbia University, New York, NY, USA

<sup>d</sup>Vagelos College of Physicians and Surgeons, Columbia University, New York, New York, U.S.A

### Abstract

Complex interventions are ubiquitous in healthcare. A lack of computational representations and information extraction solutions for complex interventions hinders accurate and efficient evidence synthesis. In this study, we manually annotated and analyzed 3,447 intervention snippets from 261 randomized clinical trial (RCT) abstracts and developed a compositional representation for complex interventions, which captures the spatial, temporal and Boolean relations between intervention components, along with an intervention normalization pipeline that automates three tasks: (i) treatment entity extraction; (ii) intervention component relation extraction; and (iii) attribute extraction and association. 361 intervention snippets from 29 unseen abstracts were included to report on the performance of the evaluation. The average F-measure was 0.74 for treatment entity extraction on an exact match and 0.82 for attribute extraction. The F-measure for relation extraction of multi-component complex interventions was 0.90. 93% of extracted attributes were correctly attributed to corresponding treatment entities.

### Keywords:

Knowledge representation, natural language processing, evidence-based medicine, complex intervention

### Introduction

Interventions can be particularly complicated due to the adoption of increasingly complex treatment designs in healthcare. According to the complex interventions working group, complex interventions combine multiple components [1], such as mental health services and primary care [2]. For instance, a study from the Effective Health Care Program emphasized the importance of others factors beyond medication intervention, such as lifestyle, diet and exercise, for preventing primary breast cancer [3].

The Participant, Intervention, Comparator, Outcome (PICO) framework is a well-adopted framework for formulating clinical queries in support of evidence-based medicine (EBM) [4]. A few systems have been developed to automatically extract PICO elements from scientific literature [5-7]. The current state-of-the-art solution for PICO information extraction is Trialstreamer [6], which supports the query of an individual intervention term (e.g., “remdesivir”, “standard of care”) but misses some valuable information, such as the strength of a medication or sophisticated treatment combinations. Even though the existing systems may map PICO elements to concepts in standard vocabularies (e.g.,

MeSH terms [7] and UMLS concepts [5]), they fail to comprehensively normalize PICO elements into compositional representations.

Previous works on intervention normalization, MedXN [8] and MedEx [9], extracted medications and corresponding attributes (e.g., strength, route and frequency) from clinical narratives and generated structured output; however, they only focused on medications. As stated in the Clinicaltrials.gov glossary [10], interventions include drugs, medical devices, procedures, vaccines and noninvasive therapies. Moreover, existing normalization systems merely extract individual medications without preserving the relations among multiple treatments in the context of combination therapies. HemOnc [11] proposed representations for combination therapies; however, its scope is limited to chemotherapy regimens.

In this paper, we present a novel representation for complex interventions and an information extraction pipeline to automatically parse free-text complex interventions from RCT abstracts into a structured representation. Our contributions are two-fold: first, we designed a compositional representation to associate all intervention components with their relations and attributes; and second, we developed an automated pipeline to normalize complex interventions into the computable representation to enable structured queries. This pipeline extends the PICO framework to support more granular queries on complex interventions for medical evidence computing.

### Methods

#### Data

Two independent datasets were created and used in our study for the development and testing of our representation and parsing pipeline of complex interventions, respectively. The development dataset includes 3,447 manually annotated intervention snippets from 261 RCT abstracts on PubMed. The testing dataset contains 361 intervention snippets extracted from 29 unseen RCT abstracts on PubMed. The intervention snippets in both datasets were annotated by two independent annotators, TK and LS.

#### Compositional representation for complex interventions

The semantic network of this representation is shown in Figure 1. An intervention component is an atomic treatment entity represented with its negation status and modifying attributes. The categories for the treatment entity include *Drug*, *Procedure*, *Device* and *Activity*. Each treatment entity is

normalized to a UMLS concept and a negation status of “affirmed” or “negated.” The attributes include *Duration* (e.g., “8 week”, “4 days”), *Frequency* (e.g., “every 12 h”, “weekly”), *Strength* (e.g. “400 mg”, “0.1%”), *Route* (e.g., “oral”, “intravenous”), *Dosage* (e.g. 2 in “2 tablet”, “single dose”), *Form* (e.g., “ophthalmic solution”, “tablet”), *Volume* (“4 ml”) and *Count* (two in “two IMZ implants”), some of which are specific to a particular treatment entity type. For example, *Dosage* is an attribute specific to *Drug* entities while *Frequency* can be an attribute specific to *Drug*, *Procedure* or *Activity* entities. There are six types of component relations, with their details and examples shown in Table 1. A component relation is directional if the order of treatment components matters; otherwise, a relation is non-directional. *Combined\_with* and *Or* are non-directional relations, whereas the remaining four relations are all directional. A complex intervention can contain more than one relation.

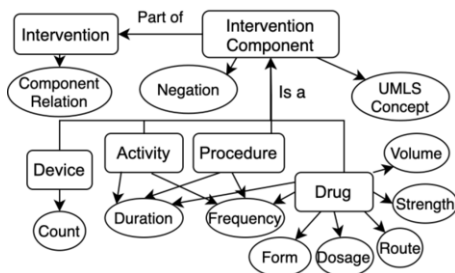


Figure 1- Semantic network of the intervention

Table 1- Types and examples of intervention relations

Relation	Example
Combined_with	<i>mixture of ropivacaine 0.1 % with clonidine 2 microg/kg, rifampicin 600 mg/day and doxycycline 200 mg/day for 45 days</i>
Or	<i>plasma or whole blood exchange, continuous hemofiltration or hemodiafiltration</i>
Before	<i>four cycles of FEC 100 then eight cycles of FEC 50, didanosine was given 1 h after delavirdine</i>
Substituted_for	<i>carteolol was substituted for timolol, switched from CFC albuterol to HFA albuterol</i>
Contain	<i>parenteral nutrition containing micronutrients, Beverage Containing Dispersible Yeast β-Glucan</i>
Without	<i>radiation without chemotherapy, physiotherapy treatment without active training</i>

**Overview of intervention normalization pipeline**

We developed a pipeline to automatically convert the free-text of a complex intervention to a structured output. As shown in Figure 2, two intervention components were extracted from the narrative “brodalumab 210 mg every 2 weeks after receiving ustekinumab through 52 weeks” – the Drug “brodalumab” with the *Strength* attribute “210 mg” and a *Frequency* of “every 2 weeks” and the Drug “ustekinumab” with the *Duration* attribute “52 weeks”. The relation between the two intervention components is *Before* (“ustekinumab” *Before* “brodalumab”).

We utilized an off-the shelf biomedical concept extraction tool - QuickUMLS [12] with part-of-speech (POS) tagging based post-coordination for treatment entity extraction and concept normalization. We then developed a pattern-based method for extracting component relations as well as component attributes. A hybrid rule-based and dependency parser-based method was developed for attribute association. A well-known biomedical NLP toolkit – scispaCy [13] was used to obtain the POS tagging and dependency parsing results of a given intervention snippet. Method details are provided below.

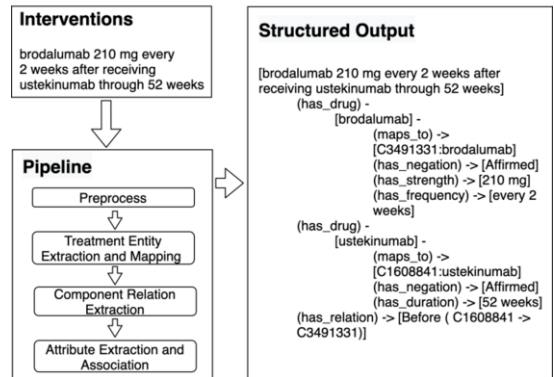


Figure 2- Intervention normalization pipeline

**Step 1. Preprocessing**

The preprocessing steps include abbreviation resolution and basic text cleaning. An abbreviation is expanded if its full name is given. For example, “SOC” is expanded to “standard of care” if “standard of care (SOC)” appears in the abstract. The abbreviation resolution is based on the algorithm proposed by Schwartz & Hearst [14] in scispaCy. To clean the text, we removed special characters such as parentheses and hyphens.

**Step 2. Treatment entity extraction and mapping**

We employed QuickUMLS to extract treatment entities and mapped them to UMLS concepts. We focused our extraction on 34 semantic types that belong to *Chemicals & Drugs*, *Activities*, *Procedures* and *Device* semantic groups defined by UMLS. For treatment entities with additional descriptive terms, we expanded the boundary of the initially extracted entities by applying a POS tagging-based method for post-coordination. For example, if the POS tag of an entity’s left adjacent token is “ADJ”, this token is merged into the entity. Other selected POS tags are “NOUN”, “ADV”, “PROPN”, “PRON” and “X”. This iterative process repeats until the beginning of the snippet is reached or the POS tag is not within the selected list.

After this step, terms are expanded. For instance, “exercise” may be expanded to “supervised trunk exercise”, “radiation” may be expanded to “intraoperative beta radiation” and “cutting” may be expanded to “hard tissue cutting”. We also expanded the boundary if an adjacent term is another extracted treatment term; for example, adjacent terms such as “povidone iodine” and “sinus irrigations” are merged to become “povidone iodine sinus irrigations”. A list of terms were manually selected to exclude from concept mapping for treatment attributes such as “intravenous” and “injection”, which are the route and form of a drug, or uncommon synonyms of a treatment concept (e.g., “control” is the synonym of the concept “C0728976: Control brand of phenylpropranolamine”, and “combinations” is the synonym of “C3539181: Short-acting sulfonamide combinations”).

### Step 3. Component relation extraction

The patterns of relations were examined and manually parsed using regular expressions. (e.g., *Combined\_with*: ENTITY0.\*\b(TRIGGER)\b.\*?ENTITY1) ENTITY0 and ENTITY1 represent the treatment entities, and TRIGGER is a list of trigger phrases, such as “added to” and “combined with”. The relation *Combined\_with* is n-ary, so it applies to multiple intervention components in the same snippet, such as “triple combination of interferon beta-1b, lopinavir-ritonavir, and ribavirin”. The other five relations are binary and are applicable to a pair of intervention components.

### Step 4. Attribute extraction and association

Attributes are extracted and associated with treatment entities by extending MedXN. We developed heuristic rules for attribute association. An attribute is associated with a treatment entity if it is contained within the treatment entity (e.g., *<C2912166> hydroxychloroquine sulfate <Strength> 200 MG </Strength> </C2912166>*, or it is directly adjacent to the treatment entity (e.g., “*<Strength> 2 % </Strength> <C0025384> mepivacaine </C0025384>*”), or it is connected to the treatment entity with other inner attributions (e.g., “*<Dosage> high - dose </Dosage> <Route> intravenous <Route> <C0025815> methylprednisolone </C0025815>*”). A dependency parser [15] is then used to attribute terms that are remotely connected to their corresponding treatment entities (e.g., “*<C0035608> rifampicin </C0035608> <Strength> 600 mg </Strength> / <Frequency> day </Frequency> and <C0013090> doxycycline </C0013090> <Strength> 200 mg </Strength> / <Frequency> day </Frequency> for <Duration> 45 days </Duration>*”, where the *Duration* “45 days” are remotely connected with the *Drug* “rifampicin” and the *Drug* “doxycycline”) based on the assumption that associated treatment entities have the shortest distance to the attributes. The dependency parse tree is generated by scispaCy for each intervention snippet. Nodes that contain the text belonging to the same attribute are then merged, i.e., attribute nodes. The property of the dependency parse tree is maintained during the process. We ensured a single parent node for the attribute node by locating the nearest common ancestor between two nodes with the partial attribute text, and merging all nodes along the path. The children of both nodes are then connected to the attribute node. The distances between an attribute node and each treatment entity node are calculated. We define that the distance between two directly connected nodes is 0 if the dependency relation type is “conj”, and the distance for all other relation types is “1”. The attribute is then associated with the treatment entity that has the shortest distance.

### Evaluation

Our evaluation metrics for the extraction of treatment entities, component relations and attributes are Precision (P), Recall (R) and F-measure (F). The accuracy of attribute association was also evaluated. 361 intervention snippets were used for the evaluation. Two annotators (SL and MB) independently annotated the snippets manually. The inter-annotator agreement was 89.2% with respect to intervention entity extraction, 73.3% for relation extraction and 71.9% for attribute extraction and association. The disagreement annotations were discussed and then re-annotated based on the annotators’ consensus. The structured output generated by the pipeline was then manually compared with the gold standard.

### Results

Figure 3 shows the summary of the evaluation dataset. Duplicates and snippets that contain only anaphoric references (e.g., “either group”, “control”) are removed. We then evaluated the performance based on the remaining 170 intervention snippets. The complexity of the text snippets is shown in Figure 3a. Most of the snippets (128 snippets) contain only one treatment entity, 33 snippets contain two entities and 9 snippets contain more than two entities. As for the distribution of treatment entity types shown in Figure 3b, there were a total of 83 *Drugs*, 106 *Procedures*, 19 *Devices* and 13 *Activities*. For the intervention component relations shown in Figure 3c, 43 relations were identified, where the *Combined\_with* relation appeared 32 times, the *Before* relation appeared 6 times and the *Or* and *Contain* relations appeared 3 and 2 times, respectively. For the attribute type shown in Figure 3d, *Route*, *Duration*, *Frequency* and *Strength* had 21, 19, 15 and 16 counts, respectively, while there were only 7 *Dosage* references and 6 *Form* references.

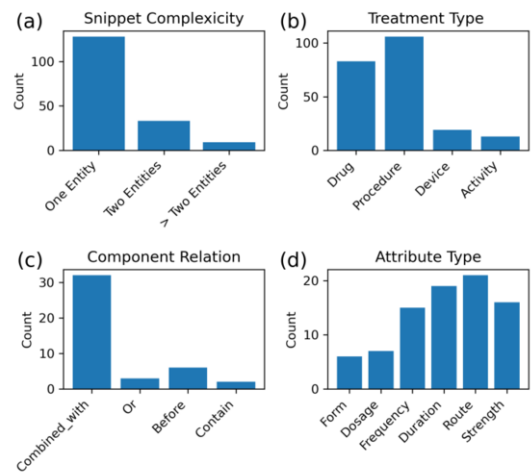


Figure 3- Statistics on the evaluation dataset

We evaluated both exact match and partial match on the extracted treatment entities. Exact match is defined as an exact matching of the string boundaries, while partial match is defined as a partial matching of the string boundaries. For the partial matching, if two entities are both partially matched to one gold-standard entity, then they are regarded as one match. Table 2 shows the evaluation results on the treatment entity extraction by reporting scores of P, R and F.

As shown in Table 2, *Procedure*, *Drug* and *Activity* reached high F-measures, while *Device* had a relatively low F-measures (0.41) on exact match. There are several possible reasons for the low F-measure for *Device*, including: (i) some terms such as “instruments” cannot map to a proper device concept on UMLS; (ii) it is difficult to identify the boundary of some device terms that contain numbers or punctuations (e.g., “q-switched nd:yag 532-nm laser”); and (iii) the current evaluation dataset has a limited sample size on *Device* entities.

Table 2- Evaluation on intervention treatment extraction

	N	Exact Match			Partial Match		
		P	R	F	P	R	F
Drug	83	0.85	0.82	0.83	0.97	0.93	0.95
Activity	13	0.90	0.95	0.93	0.90	0.95	0.93
Procedure	106	0.80	0.80	0.80	0.94	0.93	0.94
Device	19	0.46	0.38	0.41	0.92	0.75	0.83

Table 3 shows the evaluation results for the attribute extraction. *Frequency*, *Route*, *Dosage* and *Strength* have high F-measure values, while the F-measure for *Duration* is only 0.59. The low F-measure for *Duration* is mainly due to its various descriptions (e.g., “*ustekinumab* <*Duration*> *through week 52* </*Duration*>” and “<*Duration*> *five annual rounds* </*Duration*> of LDCT screening.”).

Table 3-Evaluation on attribute extraction

	N	P	R	F
Frequency	15	1.00	0.71	0.83
Route	21	1.00	0.69	0.82
Duration	19	0.73	0.50	0.59
Dosage	8	0.71	1.00	0.83
Strength	15	1.00	0.76	0.87
Form	6	1.00	1.00	1.00

To evaluate the intervention component relation extraction, we first removed snippets that have component relations from the gold standard but contain errors in the treatment entity extraction (N=8) or nested relations that are beyond the current scope (N=3, e.g., “*IFR program for 2 weeks in addition to standard care after surgery*”). As shown in Table 4, 23 relations were successfully extracted, two *Combined\_with* relations were missed and three *Combined\_with* relations were falsely extracted. The precision for relation extraction was 0.88, recall was 0.92 and F-measure was 0.90. For attribute association, 93% of the attributes were correctly associated with the corresponding treatment entities.

Table 4-Evaluation on component relation extraction

	Correctly Extracted	Incorrectly Extracted	Missed
Combined_with	18	3	2
Contain	2	0	0
Before	2	0	0
Or	1	0	0

## Discussion

We developed a representation for complex interventions in RCT publications along with a normalization pipeline to automatically structure and standardize textual descriptions of complex interventions. This pipeline promises to support more granular or in-depth queries over medical evidence containing complex interventions.

Compared with similar systems such as MedXN and MedEx, our system has two major advantages. First, we expanded the intervention types from *Drug* to include *Drug*, *Procedure*, *Activity* and *Device*. As shown in Figure 3, *Drug* interventions only constituted 38% of all the identified interventions. Another advantage of our solution is its superior performance on optimizing boundary detection for long and complex treatment entities. Our two-step mapping strategy contributed to this improvement. The first step aims to extract the terms that have a high score mapping to UMLS treatment concepts and the second step applies syntactic-based rules that work well to expand the extracted terms.

Another advantage is the introduction of the component relations. Six types of relations can be recognized with a list of patterns and trigger words after exhausting all snippets in the training dataset. Those relations could provide temporal

(*Before*, *Substituted\_for*), spatial (*Contain*) and Boolean (*Combined\_with*, *Or*, *Without*) type information between the intervention components.

For the attribute association, we proposed a hybrid approach integrating dependency parsing and pattern matching. A detailed analysis of the evaluation dataset showed that our rules cover most of the association cases. For example, the hybrid method successfully resolved cases such as “*hyperbaric 1.5 % lidocaine with 10 microg of fentanyl*”, where the heuristic rules link the *Strength* “1.5%” to the *Drug* “*lidocaine*” and the dependency-parsing based method links the *Strength* “10 microg” to the *Drug* “*fentanyl*”. Compared with the text window method proposed by MedXN, our method supports cases where an attribute maps to multiple entities. For example, “*short-term combination of an angiotensin-converting enzyme inhibitor and an angiotensin receptor blocker*” could have its “short-term” *Duration* associated with both *Drug* “*angiotensin-converting enzyme inhibitor*” and *Drug* “*angiotensin receptor blocker*.”

## Limitations and future work

Ambiguous terms contribute greatly to the inaccuracies in extraction. Typically, those terms could be the name of multiple concepts. For example, “medium” could be a quantitative concept (e.g., C0439536: Medium) or a reagent (e.g., C0010454: Culture Media) that belongs to *Chemicals & Drugs* semantic group. Therefore, a list of terms was excluded based on an examination of all the snippets in the training dataset during the development phase. Those terms include “controls”, “matched”, “men”, “group”, “baseline”, “change” and “session.” In the future, we will implement disambiguation algorithms [16] to improve coverage.

Another challenge complicating entity extraction is coordination ellipsis [17]. It exists both in the treatment entities (e.g., “*uniportal and multiportal video-assisted thoracoscopic surgery, spinal and general anesthesia*”) and intervention attributes (e.g., “*iodixanol or iohexol containing either 9 or 12 g of iodine, plain mepivacaine in a 60- or 80-mg dose*”). Sophisticated methods are needed to parse coordination ellipsis. In the future, we will test a graph-based method to reconstruct the entities [18].

Challenges also exist in the attribute association. For example, in the intervention snippet “*lopinavir 400 mg and ritonavir 100 mg every 12 h*”, the *Frequency* “every 12 h” should be associated with both *Drug* “*lopinavir*” and “*ritonavir*”, while our heuristic rule only associate it with the *Drug* “*ritonavir*”. However, in the snippet “*tazobactam 1.5 g/day plus tobramycin 3 mg/kg per day*”, the *Frequency* “per day” is correctly associated with the *Drug* “*tobramycin*”. A more refined rule is required to deal with such cases. Besides, there may exist sequential changes in the attributes, e.g., “*fluoxetine 20 mg/day for a week which was later increased to 40 mg/day*”. Currently, those attributes are extracted without capturing how they are managed. In the future, we will modify our model and pipeline to incorporate this information.

For the component relation extraction, the three incorrectly extracted relations are all *Combined\_with* and belong to two cases: (i) “*praliquat, an oral soluble guanylate cyclase stimulator*”, where “*an oral soluble guanylate cyclase stimulator*” is a description of the previous entity “*praliquat*”; and (ii) “*between the magnesium and placebo*”, where the entity “*magnesium*” and “*placebo*” are two individual interventions instead of a combined intervention. Other rules are required to filter those misidentified cases. The false negative relations in the evaluation dataset were due to the fact that the trigger words were not covered in the development phase. For example, the

phrase “experimental group (n=30) used oral cryotherapy during the infusion of the 5-FU antineoplastic agent” involves the word “during,” which could be the trigger for the relation *Combined\_with*. In the future, we plan to utilize the shortest dependency path [15] to help identify possible trigger words, explore other possible relations and extract the nested relations.

## Conclusions

We developed and validated an intervention representation along with an automated normalization pipeline for complex interventions in PubMed abstracts, aiming to convert free-text RCT intervention descriptions to a structured output represented as standard concepts. This representation involves intervention components and associated intervention attributes, as well as temporal, spatial and Boolean type relations between the intervention components. Our work also improved the extraction accuracy of highly granular treatment entities and treatment-attribute associations.

## Acknowledgments

This work was supported by 5R01LM009886-11 (Bridging the semantic gap between research eligibility criteria and clinical data; PI: Weng).

## References

- [1] M.P. Kelly, J. Noyes, R.L. Kane, C. Chang, S. Uhl, K.A. Robinson, S. Springs, M.E. Butler, and J.M. Guise, AHRQ series on complex intervention systematic reviews—paper 2: defining complexity, formulating scope, and questions, *J Clin Epidemiol* **90** (2017), 11-18.
- [2] M. Butler, R.L. Kane, D. McAlpine, R. Kathol, S.S. Fu, H. Hagedorn, and T. Wilt, Does integrated care improve treatment for depression? A systematic review, *J Ambul Care Manage* **34** (2011), 113-125.
- [3] J.-M. Guise, E.K. Nakamoto, L. LaBrant, R. Relevo, C. O’Haire, C. Most, and E. Graham, Future Research Needs To Reduce the Risk of Primary Breast Cancer in Women: Identification of Future Research Needs from Comparative Effectiveness Review No. 17, (2011).
- [4] X. Huang, J. Lin, and D. Demner-Fushman, Evaluation of PICO as a knowledge representation for clinical questions, in: *AMIA annual symposium proceedings*, American Medical Informatics Association, 2006, p. 359.
- [5] T. Kang, S. Zou, and C. Weng, Pretraining to recognize PICO elements from randomized controlled trial literature, *Studies in health technology and informatics* **264** (2019), 188.
- [6] I.J. Marshall, B. Nye, J. Kuiper, A. Noel-Storr, R. Marshall, R. Maclean, F. Soboczenski, A. Nenkova, J. Thomas, and B.C. Wallace, Trialstreamer: A living, automatically updated database of clinical trial reports, *Journal of the American Medical Informatics Association* **27** (2020), 1903-1912.
- [7] C. Schardt, M.B. Adams, T. Owens, S. Keitz, and P. Fontelo, Utilization of the PICO framework to improve searching PubMed for clinical questions, *BMC medical informatics and decision making* **7** (2007), 1-6.
- [8] S. Sohn, C. Clark, S.R. Halgrim, S.P. Murphy, C.G. Chute, and H. Liu, MedXN: an open source medication extraction and normalization tool for clinical text, *Journal of the American Medical Informatics Association* **21** (2014), 858-865.
- [9] H. Xu, S.P. Stenner, S. Doan, K.B. Johnson, L.R. Waitman, and J.C. Denny, MedEx: a medication information extraction system for clinical narratives, *Journal of the American Medical Informatics Association* **17** (2010), 19-24.
- [10] ClinicalTrials.gov, (2020). <https://clinicaltrials.gov/ct2/about-studies/glossary> (accessed April 26, 2021).
- [11] J.L. Warner, D. Dymshyts, C.G. Reich, M.J. Gurley, H. Hochheiser, Z.H. Moldwin, R. Belenkaya, A.E. Williams, and P.C. Yang, HemOnc: A new standard vocabulary for chemotherapy regimen representation in the OMOP common data model, *Journal of biomedical informatics* **96** (2019), 103239.
- [12] L. Soldaini and N. Goharian, Quickmuls: a fast, unsupervised approach for medical concept extraction, in: *MedIR workshop, sigir*, 2016, pp. 1-4.
- [13] M. Neumann, D. King, I. Beltagy, and W. Ammar, Scispacy: Fast and robust models for biomedical natural language processing, *arXiv preprint arXiv:1902.07669* (2019).
- [14] A.S. Schwartz and M.A. Hearst, A simple algorithm for identifying abbreviation definitions in biomedical text, in: *Biocomputing 2003*, World Scientific, 2002, pp. 451-462.
- [15] K. Fundel, R. Küffner, and R. Zimmer, RelEx—Relation extraction using dependency parse trees, *Bioinformatics* **23** (2007), 365-371.
- [16] Y. Wang, M. Wang, and H. Fujita, Word sense disambiguation: A comprehensive knowledge exploitation framework, *Knowledge-Based Systems* **190** (2020), 105030.
- [17] J. Beavers and I.A. Sag, Coordinate ellipsis and apparent non-constituent coordination, in: *The Proceedings of the 11th International Conference on Head-Driven Phrase Structure Grammar*, 2004, pp. 48-69.
- [18] C. Yuan, Y. Wang, N. Shang, Z. Li, R. Zhao, and C. Weng, A graph-based method for reconstructing entities from coordination ellipsis in medical text, *Journal of the American Medical Informatics Association* **27** (2020), 1364-1373.

## Address for correspondence

Chunhua Weng, chunhua@columbia.edu. Department of Biomedical Informatics, Columbia University. PH-20, 622 W 168 ST, PH-20, New York, NY 10032, USA.