

An Automated Approach for Identifying Erroneous IS-A Relations in SNOMED CT

Ran Hu, MS¹, Jay Shi, MD², Licong Cui, PhD^{1,*}, Rashmie Abeysinghe, PhD^{3,*}

¹McWilliams School of Biomedical Informatics, The University of Texas Health Science Center at Houston, Houston, TX

²Intermountain Healthcare, Denver, CO

³Department of Neurology, McGovern Medical School, The University of Texas Health Science Center at Houston, Houston, TX

Abstract

SNOMED CT is the most comprehensive clinical terminology employed worldwide and enhancing its accuracy is of utmost importance. In this work, we introduce an automated approach to identifying erroneous IS-A relations in SNOMED CT. We first extract linked concept-pairs from which we generate Term Difference Pairs (TDPs) that contain differences between the concepts. Given a TDP, if the reversed TDP also exists and the number of linked-pairs generating this TDP is less than those generating the reversed TDP, then we suggest the former linked-pairs as potentially erroneous IS-A relations. We applied this approach to the Clinical finding and Procedure subhierarchies of the 2022 March US Edition of SNOMED CT, and obtained 52 potentially erroneous IS-A relations and a candidate list of 48 linked-pairs. A domain expert confirmed 41 out of 52 (78.8%) are valid and identified 26 erroneous IS-A relations out of 48 linked-pairs demonstrating the effectiveness of the approach.

1. Introduction

Biomedical terminologies provide precise, formal descriptions of concepts and their relations in biomedicine. Biomedical applications, such as decision support systems and natural language processing (NLP) applications, use biomedical terminologies as knowledge sources¹. Quality assurance is an essential part of the management of biomedical terminologies. The goal of quality assurance is to evaluate and improve the quality of terminologies by enhancing their consistency, coverage, and completeness, eliminating redundancy, and ensuring clarity². Most biomedical terminologies have internal quality assurance mechanisms and take user feedback as a source to improve their coverage and consistency further. However, it is not practical for such approaches to cover every possible error. Some widely used biomedical terminologies are extremely large in terms of the number of concepts they contain and are updated regularly. Therefore, it is not feasible to manually review these terminologies to identify and remediate the errors. Hence, there is a pressing need for automated or semi-automated methods that can uncover errors that require significantly less manual effort³.

The Systematized Medical Nomenclature for Medicine–Clinical Terminology (SNOMED CT) is the world’s largest and most comprehensive clinical terminology, which has been utilized on a global scale to enhance the quality of medical care and the safety of patients, as well as to promote research and public health reporting⁴. SNOMED CT provides a standardized way for coding and communicating clinical concepts, which enables the exchange of accurate and consistent data between healthcare systems and applications. It enables the consistent representation of clinical information in electronic health records and has been mapped to other international standards. The SNOMED CT International Edition, which is monthly released, has close to 355,000 biomedical concepts in its March 2022 release⁵. The US Edition of the SNOMED CT merges the contents of the US Extension and the International release of SNOMED CT⁶. Each new release contains new concepts added to SNOMED CT due to the evolving needs of healthcare and the contributions of the healthcare community. The 2022 March release of the US Edition of SNOMED CT has over 360,000 active concepts⁷. The semantic annotations in SNOMED CT are machine-readable and can be used to develop models and applications in healthcare^{8,9}. SNOMED CT concepts are modeled under 19 different domains or subhierarchies, as shown in Figure 1. “Clinical Finding” and “Procedure” are two of the largest subhierarchies in SNOMED CT.

Numerous computational methods have been previously investigated for quality assurance of biomedical terminolo-

*Corresponding authors. Email: licong.cui@uth.tmc.edu; rashmie.abeyasinghe@uth.tmc.edu




Figure 1: Subhierarchies of SNOMED CT as well as the number of descendants each subhierarchy has.

gies¹⁰. The study by Burse et al. investigated the feasibility of utilizing stopwords in biological concept names as a means of completing the logical definitions of concepts that are not specified adequately by identifying lexical and semantic patterns in SNOMED CT¹¹. Gu et al. developed a semi-automated method that utilizes the SNOMED CT hierarchies to validate the semantic type assignments¹². Ochs et al. proposed a subject-based abstraction network method for quality assurance of large hierarchies in SNOMED CT¹³. They also proposed a tribal abstraction network method for auditing SNOMED CT¹⁴. The article by Abeysinghe et al. introduced an automated method for auditing inconsistency in IS-A (child-parent) relations in GO³. For SNOMED CT and National Cancer Institute (NCI) Thesaurus, various studies have shown to be effective in identifying missing IS-A relations using lexical patterns^{1, 15, 16}. Though it has not been widely employed, deep learning has also been recently investigated for missing IS-A identification^{17, 18}.

Erroneous IS-A relations are exceptionally rare when compared to missing IS-A relations. Their impact differs from missing IS-A relations in a fundamental way, as missing IS-A relations signify the absence of information, whereas erroneous IS-A relations suggest that the information is incorrect or inaccurate. Erroneous IS-A relations in SNOMED CT will negatively affect the outcome of its downstream applications. For example, predictive models and clinical decision support systems based on data involving erroneous IS-A relations in SNOMED CT might be less accurate.

Some of our earlier work in identifying missing IS-A relations in various biomedical terminologies also uncovered some erroneous IS-A relations as a byproduct of manual evaluation by domain experts^{3, 15, 19, 20}. However, it is a time-consuming and labor-intensive task to uncover such errors manually. To our knowledge, this work is the first systematic and automated approach that specifically targets identifying erroneous IS-A relations in SNOMED CT.

2. Method

In this study, we employ an automated technique to identify potential erroneous IS-A relations in the 2022 March US Edition of SNOMED CT. First, we extract the active concepts and active IS-A relations from the SNOMED CT release

files. Then, we extract all the hierarchically linked concept-pairs. From these linked concept-pairs, we further extract Term Difference Pairs (TDPs) that describe the differences between the linked concept-pair. Then, we analyze all such TDPs to see whether, for a given TDP, a reversed TDP exists. Afterward, the numbers of linked concept-pairs that derived the TDP and the reversed TDP are counted. If the TDP is exhibited by less number of linked concept-pairs than the reversed TDP, then these concept-pairs are deemed to indicate erroneous IS-A relations or vice versa.

2.1 Generating linked concept-pairs

After extracting all active concepts and active IS-A relations from SNOMED CT release files, we leverage the python package NetworkX²¹ to represent SNOMED CT as a graph. With NetworkX, we obtain all the direct and indirect hierarchically linked concept-pairs in SNOMED CT.

A linked concept-pair (A, B) is generated from a concept-pair A and B that have an IS-A relation such that: A IS-A B . Note that linked concept-pairs are directional, i.e., $(A, B) \neq (B, A)$. A more specific example would be the concept 866195005: *Ultrasonography of joint of upper limb (procedure)* IS-A concept 449921006: *Ultrasonography of upper limb (procedure)*. This cannot be the other way around. Generally, both (A, B) and (B, A) would not exist in the terminology as the relations A IS-A B and B IS-A A will form a cycle.

Linked concept-pairs are generated from indirect IS-A relations as well. For instance, if SNOMED CT records IS-A relations A IS-A B and B IS-A C , then we generate linked concept-pairs (A, B) ; (B, C) ; and (A, C) . Note that the IS-A relation between A and C is indirect (or transitive). Figure 2 shows an example of linked concept-pair extraction based on the IS-A relations in SNOMED CT.




Figure 2: The concepts 866195005: *Ultrasonography of joint of upper limb (procedure)*, 449921006: *Ultrasonography of upper limb (procedure)*, and 710340004: *Ultrasonography of limb (procedure)* generates three linked concept-pairs: (866195005, 449921006); (449921006, 710340004); and (866195005, 710340004).

2.2 Obtaining lexical features for SNOMED CT concepts

Each SNOMED CT concept has a Fully Specified Name (FSN) that unambiguously explains the meaning of the concept. For each concept C , we obtain a set of lexical features $L(C)$ from its FSN by following the steps below in the given order:

1. Convert the FSN to lowercase;
2. Break down the FSN into words; and
3. Eliminate duplicated words and consider them as an unordered set of words.

For example, the concept with the identifier 403914000 has the FSN *Superficial basal cell carcinoma (disorder)*. The lexical features obtained from this concept would be $\{ \text{'superficial'}, \text{'basal'}, \text{'cell'}, \text{'carcinoma'} \}$.

2.3 Obtaining Term Difference Pairs (TDPs)

A term difference pair $TDP(A, B)$ could be obtained by a linked concept-pair (A, B) and the lexical features $L(A)$ and $L(B)$ as follows:

$$TDP(A, B) = [X, Y]$$

where $X = L(A) - L(B)$ and $Y = L(B) - L(A)$.

Essentially, the TDP accentuates the differences between the lexical features of the linked concept-pair. For instance, the concepts *Superficial basal cell carcinoma (disorder)* with lexical features {'superficial', 'basal', 'cell', 'carcinoma'} and *Basal cell carcinoma of skin (disorder)* with lexical features {'basal', 'cell', 'carcinoma', 'of', 'skin'} will generate the TDP [{'superficial'}, {'of', 'skin'}]. Note that the common lexical features {'basal', 'cell', 'carcinoma'} have been removed from both concepts' lexical features to generate the TDP.

The TDP is directional, i.e., $TDP(A, B) \neq TDP(B, A)$. If both the sets in a TDP are empty sets, then we ignore such TDPs. In addition, if both the sets in a TDP contain only the stop words 'with', 'of', 'and', 'or', 'and/or', 'the', 'AND/OR', 'no', 'not', 'without', 'due to', 'secondary to', 'except', 'by', 'after', 'able', 'removal', 'replacement', 'NOS', then such TDPs are also ignored. This is because the TDP is supposed to highlight important lexical differences between the concept-pairs, and if it only contains stop words, then this purpose is not satisfied.

Note that the same TDP could be generated by multiple concept-pairs. For example, Figure 3 shows the TDP [{'skin', 'of'}, {'superficial'}] obtained by linked concept pair *Foreign body in skin of lip (disorder)* and *Superficial foreign body in lip (disorder)*. The same TDP [{'skin', 'of'}, {'superficial'}] can also be generated by linked concept-pair *Puncture wound of skin (disorder)* and *Superficial puncture wound (disorder)*. Once TDPs generated by all linked concept-pairs are identified, the concept-pairs that generate the same TDP are grouped together.




Figure 3: The TDP generated from concept 298065007: *Foreign body in skin of lip (disorder)* and concept 874837003: *Superficial foreign body in lip (disorder)*.

2.4 Identifying potentially erroneous IS-A relations

Let $TDP_1 = [C, D]$ be generated by m linked concept-pairs and $TDP_2 = [D, C]$ be generated by n linked concept-pairs. If $m > n$, then we suggest that all the n linked concept-pairs that generate TDP_2 represent potentially erroneous IS-A relations. In other words, if two TDPs are found to be the reverse of each other, then the one that is generated by a less number of linked concept-pairs is considered to be denoting potentially erroneous IS-A relations, and hence all the linked concept-pairs that generated that particular TDP is considered to be potentially erroneous.

For example, in Figure 4, the TDP [{'superficial'}, {'skin', 'of'}] is obtained by linked concept-pair *Superficial basal cell carcinoma (disorder)* and *Basal cell carcinoma of skin (disorder)*. The TDP [{'skin', 'of'}, {'superficial'}] is obtained by linked concept-pair *Foreign body in skin of lip (disorder)* and *Superficial foreign body in lip (disorder)*. Two additional linked concept-pairs *Puncture wound of skin (disorder)* and *Superficial puncture wound (disorder)*, and *Foreign body in skin of head (disorder)* and *Superficial foreign body in head (disorder)* also generate the same TDP [{'skin', 'of'}, {'superficial'}]. The count of TDP [{'superficial'}, {'skin', 'of'}] versus the count of TDP [{'skin', 'of'}, {'superficial'}] is 1:3. In this instance, the TDP [{'superficial'}, {'skin', 'of'}] is suggested to indicate potentially erroneous IS-A relations and hence all the linked concept-pairs that generated this TDP including *Superficial basal cell carcinoma (disorder)* and *Basal cell carcinoma of skin (disorder)* is considered to be potentially erroneous.




Figure 4: The TDP $\{\{\text{'superficial'}\}, \{\text{'skin'}, \text{'of'}\}\}$ obtained by linked concept-pair *Superficial basal cell carcinoma (disorder)* and *Basal cell carcinoma of skin (disorder)* is the reverse of the TDP $\{\{\text{'skin'}, \text{'of'}\}, \{\text{'superficial'}\}\}$ obtained by linked concept-pair *Foreign body in skin of lip (disorder)* and *Superficial foreign body in lip (disorder)*.

Note that we discard linked concept pairs with $\text{TDP} = [S, \emptyset]$ (where $S \neq \emptyset$) from the results obtained above. This is because in a TDP, if the second set is empty, then in the corresponding linked concept pair, the lexical features of the right concept are a subset of that of the left concept. In such cases, it is more likely that the linked concept-pair is valid. The linked concept-pair *Breast care procedure (procedure)* and *Breast procedure (procedure)* with $\text{TDP} \{\{\text{'care'}\}, \emptyset\}$ is an instance in this case.

2.5 Evaluating the identified potentially erroneous IS-A relations

We focused on the largest subhierarchies of SNOMED CT (i.e., “*Clinical Finding*” and “*Procedure*”) to evaluate the performance of our approach in identifying erroneous IS-A relations. The potentially erroneous IS-A relations obtained from these subhierarchies were evaluated by a domain expert (author JS). We also provided the domain expert with a list of linked concept-pairs in tied situations (i.e., $m = n$) to determine if there are valid cases of erroneous IS-A relations.

3. Results

Table 1: Summary statistics for the *Clinical Finding* and *Procedure* subhierarchies.

	Clinical Finding	Procedure
Number of concepts	120,609	59,489
Number of linked concept-pairs	2,522,326	1,142,984
Number of TDPs	2,111,186	976,829
Number of potentially erroneous IS-A (untied)	38	14
Number of tied linked concept-pairs	28	20

The 2022 March release of the US Edition of SNOMED CT contained 361,780 active concepts. Table 1 shows some

summary statistics related to the results obtained for the *Clinical Finding* and *Procedure* subhierarchies. For instance, the *Clinical Finding* subhierarchy contained 120,609 concepts and 2,522,326 linked concept-pairs. These generated 2,111,186 TDPs. In total, the method suggested 52 potentially erroneous IS-A relations in the untied group (i.e., $m > n$). In addition, there were 48 tied cases where the number of linked concept-pairs generating the TDP and generating the reversed TDP are the same (i.e., $m = n$). Our pipeline approximately took 20 minutes to come up with these results.

Table 2 shows ten examples of TDPs as well as the number of potentially erroneous IS-A relations suggested by each TDP. For instance, TDP $[\emptyset, \{\text{'analog'}\}]$ suggested 3 potentially erroneous IS-A relations; and TDP $[\emptyset, \{\text{'monocular'}\}]$ suggested 2 potentially erroneous IS-A relations.

Table 2: Ten examples of TDPs and the number of potentially erroneous IS-A relations suggested by each TDP.

TDP	Number of potentially erroneous IS-A
$[\emptyset, \{\text{'analog'}\}]$	3
$[\emptyset, \{\text{'group'}, \text{'blood'}\}]$	2
$[\emptyset, \{\text{'monocular'}\}]$	2
$[\{\text{'anorectal'}\}, \{\text{'rectal'}\}]$	2
$[\{\text{'lumbar'}\}, \{\text{'lumbosacral'}\}]$	2
$[\emptyset, \{\text{'attenuated'}, \text{'live'}\}]$	1
$[\{\text{'anorectal'}\}, \{\text{'anal'}\}]$	1
$[\{\text{'superficial'}\}, \{\text{'of'}, \text{'skin'}\}]$	1
$[\{\text{'syndrome'}\}, \{\text{'primary'}\}]$	1
$[\{\text{'pharyngeal'}\}, \{\text{'oropharyngeal'}\}]$	1

3.1 Evaluation results

The domain expert reviewed the 52 potentially erroneous IS-A relations identified by our method and provided detailed feedback. The domain expert found 41 (78.8%) cases to be denoting valid erroneous IS-A relations. Table 3 contains ten examples of such valid erroneous IS-A relations confirmed by the domain expert.

For the 48 tied cases, the domain expert validated 26 erroneous IS-A relations (see Table 4 for five examples).

4. Discussion

In this study, we developed a novel automated approach for detecting potential erroneous IS-A relations in SNOMED CT. The method leverages TDPs that highlight the differences between a hierarchically linked concept-pair. The TDP is supposed to indicate important lexical features unique to each concept in a linked concept-pair that make the IS-A relation hold. As such, our argument is that a TDP and its reversed TDP should not both exist in a certain terminology as the IS-A relationship is a directional relationship. Therefore, in situations where both exist, the TDP that is generated by a less number of IS-A relations is likely to be erroneous as it is the less common one.

However, it should be mentioned that in tied situations, the number of linked concept-pairs generating the TDP and the number of other linked concept-pairs generating the reversed TDP is the same. In such situations, although our method was not able to automatically suggest which ones are potentially erroneous, it is still valuable to narrow down the potential candidates.

Table 3: Ten examples of valid erroneous IS-A relations identified by our method and verified by the domain expert.

Subconcept	Superconcept
Carcinoma in situ of uterine cervix (disorder)	Squamous cell carcinoma in situ of uterine cervix (disorder)
Dissection of cerebral artery (disorder)	Dissecting aneurysm of cerebral artery (disorder)
Metabolic encephalopathy (disorder)	Toxic metabolic encephalopathy (disorder)
Suicidal (finding)	Suicidal thoughts (finding)
Anorectal fistula (disorder)	Fistula of anus (disorder)
Pelvis juvenile osteochondropathy (disorder)	Hip juvenile osteochondropathy (disorder)
Angiography of arteries of bilateral extremities (procedure)	Angiography of arteries of bilateral upper extremities (procedure)
Adenocarcinoma of pancreatic duct (disorder)	Primary adenocarcinoma of pancreatic duct (disorder)
Focal juxtapapillary choroiditis (disorder)	Juxtapapillary focal chorioretinitis (disorder)
Pethidine overdose (disorder)	Pethidine analog overdose (disorder)

Table 4: Five examples of valid erroneous IS-A relations for tied situations.

Subconcept	Superconcept
Pain of joint of left lower leg (finding)	Pain of left knee joint (finding)
Reconstruction of mandible (procedure)	Osteoplasty of mandible (procedure)
Lumbosacral agenesis (disorder)	Sacral agenesis (disorder)
Tracheobronchial hemorrhage (disorder)	Bronchial hemorrhage (disorder)
Malignant neoplasm of middle lobe bronchus (disorder)	Neoplasm of bronchus of right middle lobe (disorder)

As it leverages the concept names and hierarchical relations that most biomedical terminologies contain, our approach can be easily adopted to any other terminology to identify potentially erroneous IS-A relations. Furthermore, a TDP can be generated by a pair of concepts linked by any type of relationship. Therefore, our approach can readily be applied to audit other types of relationships as well.

Out of the erroneous IS-A relations confirmed by the domain expert, 3 were found to be reversed in the newer September 2023 US Edition of SNOMED CT. In addition, in 8 validated erroneous IS-A relations, at least one concept was found to be no longer active in the newer SNOMED CT release. For instance, the IS-A relation between the concepts *Adenocarcinoma of pancreatic duct (disorder)* and *Primary adenocarcinoma of pancreatic duct (disorder)*, that existed in the March 2022 release, has been reversed in the 2023 release.

It should also be mentioned that ontology quality assurance approaches such as this are meant to discover defects of an ontology that were previously unknown. Different approaches will discover part of the defects but not all. Hence, there could be other erroneous IS-A relations that were not identified by this method. Therefore, it is essential to develop further algorithms to identify more erroneous relationships that were not uncovered by this method.

4.1 Comparison with similar work

We have previously leveraged TDP-based approaches to audit different biomedical terminologies, albeit for a different task in identifying missing IS-A relations^{3, 19, 20, 22, 23}. The general approach was to check whether the same ITP could be found among hierarchically linked and unlinked concept-pairs. In such situations, we suggested a potentially

missing IS-A relation between the unlinked concept-pair. However, if the unlinked concept-pair was found not to form a valid IS-A relation as determined by a domain expert, then the corresponding linked concept-pair that generated the same TDP was further reviewed to uncover potentially erroneous IS-A relations. Therefore, while those previous works uncovered cases of erroneous IS-A relations, they were rather revealed as a byproduct of their main target in identifying missing IS-A relations. In addition, to reveal the erroneous IS-A relations, a domain expert had first to determine that a missing IS-A relation suggestion is invalid. In this work, while we still use TDPs, the approach leveraging the reverse TDPs is specifically designed to identify erroneous IS-A relations and is different from previous approaches. It must also be stated that erroneous IS-A relations occur very rarely in a biomedical terminology when compared to missing IS-A relations.

4.2 Limitations and future work

As mentioned earlier, in some situations, the TDP and the reversed TDP are both generated by the same number of linked concept-pairs. Hence, our approach was unable to automatically determine which TDP denotes erroneous IS-A relations by the majority rule. For example, the TDP [{‘reconstruction’}, {‘osteoplasty’}] is generated by exactly one linked concept-pair *Reconstruction of mandible (procedure)* and *Osteoplasty of mandible (procedure)*. Similarly, the reversed TDP [{‘osteoplasty’}, {‘reconstruction’}] is generated by exactly one linked concept-pair *Osteoplasty of facial bones (procedure)* and *Reconstruction of facial bones (procedure)*. This is a case of the tied group. In the future, we would like to investigate NLP techniques to see whether we can address such cases.

In this work, we leveraged the fully specified name of SNOMED CT concepts to obtain lexical features. We also attempted to leverage a synonym replacement approach by leveraging synonyms of concepts. However, no additional erroneous pairs were discovered. We may further enhance the approach in this work by leveraging the semantic meanings of concepts. We will also study whether the lexical features could be enriched by leveraging the concept’s ancestors’ lexical features.

In addition, we applied our approach to the two largest subhierarchies of SNOMED CT: *Clinical Finding* and *Procedure*. In the future, we expect to apply it to the rest of the subhierarchies in SNOMED CT as well as other biomedical terminologies so that the generalizability of the method can be further investigated.

Another interesting direction for future research would be to investigate whether the current state-of-the-art large language models (LLMs) (e.g., ChatGPT) could be leveraged for validating the potentially erroneous IS-A relations. Our preliminary exploration showed that different LLMs may generate varying responses for the same question regarding whether a concept is a subtype of another concept.

5. Conclusion

In this paper, we investigated an automated approach to identify erroneous IS-A relations in SNOMED CT. We first obtained the fully specified names of concepts as well as direct and indirect linked concept-pairs. Then, lexical features for each concept were extracted from the fully specified names. A TDP was obtained from each linked concept-pair highlighting different lexical features in each concept. If a TDP and its reversed TDP both exist, then we suggest that the TDP that was generated by the least number of linked concept-pairs indicates potentially erroneous IS-A relations. Applying this approach to the two largest subhierarchies: *Clinical Finding* and *Procedure* of SNOMED CT (2022 March US Edition), we identified 52 potentially erroneous IS-A relations as well as a candidate list of 48 linked concept-pairs. Domain expert evaluation revealed that 41 (78.8%) out of 52 cases were valid, and 26 erroneous IS-A relations were found among 48 linked concept-pairs, indicating the effectiveness of our approach in automatically identifying erroneous IS-A relations.

6. Acknowledgement

This work was supported by the National Science Foundation (NSF) through grant 2047001, and National Institutes of Health (NIH) through grants R01LM013335 and R01NS116287. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NSF or NIH.

References

1. Abeyasinghe R, Brooks MA, Talbert J, Licong C. Quality assurance of NCI Thesaurus by mining structural-lexical patterns. In AMIA annual symposium proceedings 2017 (Vol. 2017, p. 364). American Medical Informatics Association.
2. Cui L, Tao S, Zhang GQ. Biomedical ontology quality assurance using a big data approach. *ACM Transactions on Knowledge Discovery from Data (TKDD)*. 2016 May 24;10(4):1-28.
3. Abeyasinghe R, Hinderer EW, Moseley HN, Cui L. Auditing subtype inconsistencies among gene ontology concepts. In 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) 2017 Nov 13 (pp. 1242-1245). IEEE.
4. Chang E, Mostafa J. The use of SNOMED CT, 2013-2020: a literature review. *Journal of the American Medical Informatics Association*. 2021 Sep 1;28(9):2017-26.
5. 5-step briefing [Internet]. SNOMED International. [cited 2023Feb25]. Available from: <https://www.snomed.org/snomed-ct/five-step-briefing>
6. SNOMED CT [Internet]. U.S. National Library of Medicine. National Institutes of Health; [cited 2023Feb26]. Available from: <https://www.nlm.nih.gov/healthit/snomedct/>
7. SNOMED CT Managed Service - US edition release notes - March 2022 [Internet]. SNOMED Confluence. [cited 2023Feb28]. Available from: <https://confluence.ihtsdotools.org/display/RMT/SNOMED+CT+Managed+Service+-+US+Edition+Release+Notes+-+March+2022>
8. Peterson KJ, Liu H. Automating the transformation of free-text clinical problems into SNOMED CT expressions. *AMIA Summits on Translational Science Proceedings*. 2020;2020:497.
9. Wardle M, Spencer A. Implementation of SNOMED CT in an online clinical database. *Future Healthcare Journal*. 2017 Jun;4(2):126.
10. Zhu X, Fan JW, Baorto DM, Weng C, Cimino JJ. A review of auditing methods applied to the content of controlled biomedical terminologies. *Journal of biomedical informatics*. 2009 Jun 1;42(3):413-25.
11. Burse R, McArdle G, Bertolotto M. Targeting stopwords for quality assurance of SNOMED-CT. *International Journal of Medical Informatics*. 2022 Nov 1;167:104870.
12. Gu H, Chen Y, He Z, Halper M, Chen L. Quality assurance of UMLS semantic type assignments using SNOMED CT hierarchies. *Methods of information in medicine*. 2016;55(02):158-65.
13. Ochs C, Geller J, Perl Y, Chen Y, Xu J, Min H, Case JT, Wei Z. Scalable quality assurance for large SNOMED CT hierarchies using subject-based subtaxonomies. *Journal of the American Medical Informatics Association*. 2015 May 1;22(3):507-18.
14. Ochs C, Geller J, Perl Y, Chen Y, Agrawal A, Case JT, Hripcsak G. A tribal abstraction network for SNOMED CT target hierarchies without attribute relationships. *Journal of the American Medical Informatics Association*. 2015 May 1;22(3):628-39.
15. Zheng F, Shi J, Cui L. A lexical-based approach for exhaustive detection of missing hierarchical IS-A relations in SNOMED CT. In AMIA Annual Symposium Proceedings 2020 (Vol. 2020, p. 1392). American Medical Informatics Association.
16. Hao X, Abeyasinghe R, Shi J, Cui L. A substring replacement approach for identifying missing IS-A relations in SNOMED CT. In 2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) 2022 Dec 6 (pp. 2611-2618). IEEE.
17. Abeyasinghe R, Zheng F, Bernstam EV, Shi J, Bodenreider O, Cui L. A deep learning approach to identify missing is-a relations in SNOMED CT. *Journal of the American Medical Informatics Association*. 2022 Dec 20.
18. Liu H, Zheng L, Perl Y, Geller J, Elhanan G. Can a convolutional neural network support auditing of nci thesaurus neoplasm concepts?. In ICBO 2018 Aug 7.
19. Manuel W, Abeyasinghe R, He Y, Tao C, Cui L. Identification of missing hierarchical relations in the vaccine ontology using acquired term pairs. *Journal of Biomedical Semantics*. 2022 Aug 13;13(1):22.

20. Mohtashamian M, Abeysinghe R, Hao X, Cui L. Identifying Missing IS-A Relations in Orphanet Rare Disease Ontology. In 2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) 2022 Dec 6 (pp. 3274-3279). IEEE.
21. Hagberg A, Swart P, Chult D. Exploring network structure, dynamics, and function using NetworkX. Los Alamos National Lab.(LANL), Los Alamos, NM (United States); 2008 Jan 1.
22. Mohtashamian M, Hu R, Abeysinghe R, Hao X, Xu H, Cui L, Automated Identification of Missing IS-A Relations in the Human Phenotype Ontology. In AMIA annual symposium proceedings. Forthcoming 2023. American Medical Informatics Association.
23. Abeysinghe R, Zheng F, Hinderer EW, Moseley HN, Cui L. A lexical approach to identifying subtype inconsistencies in biomedical terminologies. In 2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) 2018 Dec 3 (pp. 1982-1989). IEEE.