QA4IE: A Quality Assurance Tool for Information Extraction

Rafael Jimenez Silva[1], Kaushik Gedela[1], Alex Marr[1], Bart Desmet[1], Carolyn Rosé[1,2], Chunxiao Zhou[1]

[1]National Institutes of Health Clinical Center, Bethesda, Maryland, USA
[2]Language Technologies Institute, Carnegie Mellon University, Pittsburgh, USA
{rafael.jimenezsilva, kaushik.gedela, alex.marr, bart.desmet, carolyn.rose, chunxiao.zhou}@nih.gov

Abstract

Quality assurance (QA) is an essential though underdeveloped part of the data annotation process. Although QA is supported to some extent in existing annotation tools, comprehensive support for QA is not standardly provided. In this paper we contribute QA4IE, a comprehensive QA tool for information extraction, which can (1) detect potential problems in text annotations in a timely manner, (2) accurately assess the quality of annotations, (3) visually display and summarize annotation discrepancies among annotation team members, (4) provide a comprehensive statistics report, and (5) support viewing of annotated documents interactively. This paper offers a competitive analysis comparing QA4IE and other popular annotation tools and demonstrates its features, usage, and effectiveness through a case study. The Python code, documentation, and demonstration video are available publicly at https://github.com/CC-RMD-EpiBio/QA4IE.

Keywords: quality assurance, evaluation, discrepancy analysis

1. Introduction

For the past 30 years, benchmark tasks have driven forward progress in the field of Natural Language Processing (NLP). As larger datasets have become the norm in the field, the number of available annotation tools has increased accordingly (Atdağ and Labatut, 2013; Beasley and Manda, 2018; Neves and Ševa, 2021). Though some benchmark tasks leverage training data with naturally occurring labels, common gold standard corpus construction requires substantial annotation effort. Despite advances in pipeline and model development, errors remain likely to occur. Therefore, there is a great need to develop tools that can help correct and improve annotated language resources. A comprehensive quality assurance tool helps guarantee the high quality of language resources and their downstream reliability for evaluation.

In this paper we present QA4IE, a novel tool that addresses an important gap existing in other current annotation tools, namely comprehensive support for quality assurance. Due to the importance of QA in data annotation, most of the existing annotation tools include facilities for error checking (Cejuela et al., 2014; Cunningham et al., 2013; de Castilho et al., 2014; Grosman et al., 2020; Grundke et al., 2016; Klie et al., 2018; Kummerfeld, 2019; LightTag, 2021; Ogren, 2006; Stenetorp et al., 2012), though possibly perfunctory in nature. In contrast, QA4IE offers the following both for annotation and for information extraction: First, it offers automatic detection of potential problems in the annotation results. The user provides formatting standards that are used by QA4IE to validate the annotations and alert the user(s) to deviations from the standard. Second, QA4IE provides measures of agreement between annotators in terms of both token-level and entity-level IRR measurements. Third, based on its comparison between annotations across annotators, QA4IE provides a visual display of the discrepancies between annotators. Fourth, QA4IE provides comprehensive summary statistics computed over annotations, including counts for each type, length distribution, and more. Fifth, users are provided with a corpus viewer that enables viewing annotations of the selected corpus interactively. In addition, a comprehensive QA report, including details of errors and warnings found, text and annotation statistics, annotation performance evaluation, and discrepancy analysis, can be exported in CSV or TXT format.

In the remainder of the paper we will expand upon the desiderata that motivated work on QA4IE. In the next section, we describe in detail the technical foundation for the development of QA4IE. After that we offer a case study illustrating the use of QA4IE. Finally, we provide a competitive analysis, comparing QA4IE to existing tools. The article concludes with a summary and a plan for future work.

2. Annotation Process Support Desiderata

Text annotation specifically refers to the marking of text to highlight structure and substance, for a variety of language tasks that may be operationalized either as classification tasks or sequence labeling tasks. Common annotation tasks include rhetorical structure, concepts and relationships between them, syntactic dependencies, and so on. Annotation tools support the annotation process, enabling far greater efficiency and accuracy than would be possible without them (Ide and Pustejovsky, 2017; Pustejovsky and Stubbs, 2012). In this section we explore the challenges in annotation efforts that render spot checking insufficient and therefore motivate the need for tools with specific affordances for cataloguing potential errors and making them easy to find.

The multi-level structure of natural language provides the opportunity for a variety of error types to occur. In the example of nested named entity recognition (NER) (Nadeau and Sekine, 2007), inconsistencies between annotators may occur in terms of annotation set names and the corresponding annotation schema, sub-entity annotations outside the parent entity annotations, impermissible entity overlap, inconsistent attributes, and so on.

What exacerbates the already difficult QA process is the demand for large scale annotated data, which requires that the annotation work be done by a team of annotators, not just a pair of annotators. In order to ensure the consistency of the annotation work, an inter-rater reliability (IRR) test is performed before the annotators are able to work independently (Artstein and Poesio, 2008). When the IRR is higher than the preset threshold and the main problems that cause inconsistency are resolved, the quality of each annotator's individual annotation work can be assured. Therefore, QA requires the provision of IRR metrics and calculations for teams of annotators.

Due to the constraints of project timeline and budget, it is not always possible to resolve every potential error. Instead, annotation teams must work together to prioritize the most critical issues to address. This requires a QA tool to summarize and categorize annotation inconsistencies and identify the top issues for the annotation team to prioritize. In order to accomplish this, issues must be considered in the context of the big picture of the corpus as a whole. Thus, the QA tool needs to provide summary statistics to support this reflection.

Validation of annotations against a pre-set standard is an incredibly important step that is largely not supported by existing tools. For example, it is typically not trivial to determine whether the original text has been modified in the annotation process, whether a sub-entity is beyond the scope of the parent entity, whether there are overlaps among entities/sub-entities that are not allowed, and so on.

Regarding IRR, although most tools provide the function of calculating IRR, they often provide either the entity-level or the token-level metrics, but not both, even though these two different levels of IRR metrics often provide complementary information and thus need to be considered together. In terms of discrepancy analysis, the presentation of annotation differences is often based on the assumption of one-to-one entity correspondence, which does not always hold (Cunningham et al., 2013). In addition, there is an absence of categorization for different types of discrepancies.

In short, there is a great need for QA4IE, with its comprehensive and powerful QA features for supporting both the quality and efficiency of the annotation process, increasing our understanding of data and annotation, and thus providing an important guarantee for high-quality modeling and evaluation. A more detailed comparison of QA4IE with existing tools is provided in the Competitive Analysis in Section 5.

3.   QA4IE System Description

QA4IE addresses all the desiderata discussed in the previous section. Information extraction includes Named Entity Recognition, Coreference Resolution, Named Entity Linking, Relation Extraction, Event Extraction, Terminology Extraction, and so on. In this paper, we use nested NER as an example to illustrate the features of QA4IE. Although the current version of QA4IE focuses on nested NER, these features can be easily applied to other information extraction tasks. QA4IE consists of five main components: Error checking and Validation, IRR, Discrepancy Analysis, Visualization, and Prioritization. In the following, we introduce each component one by one.

3.1   Error Checking and Validation

Error checking and Validation covers both the document and annotation levels. *Document-level validation* is the first and simplest part of QA, which sets the foundation for the subsequent portions. One aspect is to check whether all annotators are using the same basic data format for annotation consistently over the whole corpus. Although text may be pre-processed and distributed to each annotator in the same way before annotation, some annotation tools allow annotators to purposefully or accidentally change text during annotation. Differences in text may also occur if progress is stored locally rather than on a central server. These changes result in false positives in the discrepancy analysis. To address this, QA4IE checks if the text of the same document differs across the versions used by each of multiple annotators. In addition, QA4IE also check consistency in annotation set name across annotators working on the same annotation task.

Annotation level Validation is where annotations are checked for compliance with corresponding annotation guidelines. In order to provide customized services for different annotation tasks, QA4IE offers a structured configuration file, in which the user can provide the following information: (1) the location of the annotation files; (2) the location of the output files; (3) the type of annotation task: QA4IE currently only has the option of sequence labelling (other options, such as text classification, are under development) (4) and the character encoding type. (5) Finally, for sequence labeling, the user can specify the hierarchical structure among entities/sub-entities and which entities/sub-entities are allowed to overlap with each other. QA4IE also checks annotations under default assumptions. For example, by default, the annotations may not be allowed to go beyond the range of the text; or, the end index of an entity/sub-entity may not be allowed to be equal to or less than its starting index. Annotation offsets that occur outside of the scope of the text may seem impossible. This is true for most original annotations. But annotation results often need some post-processing, such as converting format, filtering irrelevant content, etc. If the post-processing process allows modifying the indices of annotations, this unexpected error may arise.

3.2   IRR

IRR, which is also called inter-annotator agreement (IAA), is commonly used to measure the quality of annotation work. Assuming that the annotations from each annotator have been validated, the next step is to ensure some level of consensus among annotators. Cohen's kappa coefficient (Artstein and Poesio, 2008) is commonly used to measure IRR for classification tasks because it is a chance-corrected metric. For sequence labeling, where chance agreement calculation is still an open problem, F1 score is usually used to approximate the IRR (Cunningham et al., 2013).

3.3   Discrepancy Analysis

IRR is just a measure of agreement that defines quality in terms of reproducibility across annotators. As a next step, the *Discrepancy Analysis* functionality is used to help the annotation team find out where the problems are, the reasons for the problems, and to distinguish between more and less important problems as part of their prioritization process. The discrepancy analysis in existing annotation
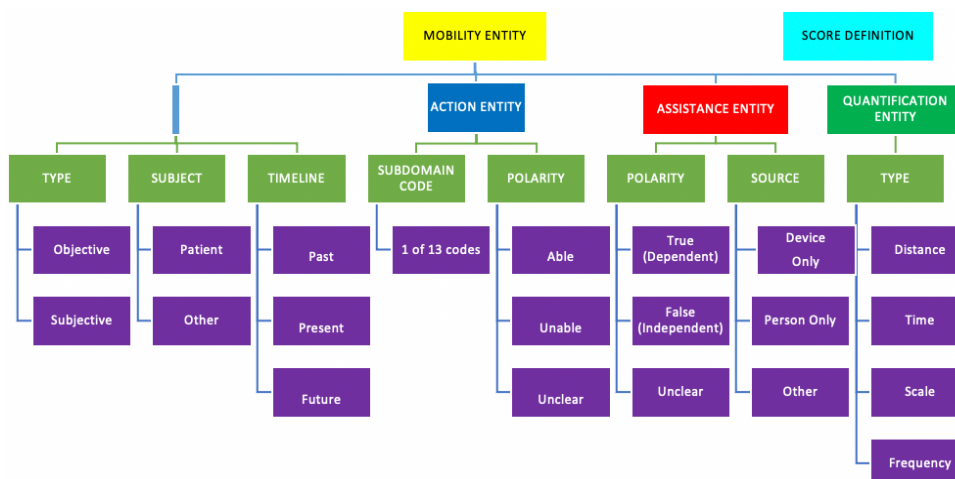
Figure 1: Mobility information structure.

software is typically only a display of differences based on the assumption of one-to-one correspondence, i.e., each unit of one annotator's annotation can be matched to a unit of another annotator's annotation if both of them annotate the corresponding unit in the same text (Cunningham et al., 2013). However, this assumption does not always hold. We find that for the annotation of long text span entities, the annotation style of each annotator is different, which often results in a many-to-many correspondence situation. This is because some annotators are accustomed to labeling a continuous long text span as several short entities, while others may be used to combining several adjacent short entities into one long entity.

3.4 Visualization

Visualization is then offered beyond simply pinpointing specific discrepancies between pairs of annotators. The display for checking one-to-one correspondences does not fully demonstrate the similarities and differences across multiple annotators. By partitioning the annotations of multiple annotators into clusters based on annotation overlap, QA4IE compares the annotations of multiple annotators within each resulting text span cluster, so that these more complex sets of differences between annotators can be seen at glance. To provide better feedback to the annotation team, QA4IE supports assignment of these clusters into the following categories: (1) not every annotator has annotated this; (2) split/combine, i.e., an entity annotated by one annotator overlaps with multiple entities of another annotator; (3) partial match; (4) corresponding entities have the same text span, but their attributes are different. QA4IE compares the discrepancies of each pair of annotators and the whole annotation team, and sorts the various types of discrepancies for the annotation team's reference.

3.5 Prioritization

Prioritization can be supported by enabling an annotation team to view identified issues in the context of the composition of the corpus. QA4IE also tallies the annotations of each annotator to better understand the characteristics of the task and the annotation style of each annotator. Summary statistics include: (1) the number of various types of entities in each document; (2) number of entities with specific attributes; (3) length of entities, and (4) the length of entities with specific attributes.

We will illustrate all features with details in the case study below.

4. Case Study

The current QA4IE is a Python based command-line tool that is intuitive and accessible. Anyone familiar with the command line interface can quickly master all of QA4IE's features. The user needs to provide the input folder for annotated GATE XML files, the output folder for QA reports, and the annotation schema specifications in a configuration file. GATE supports a range of file formats by default, including XML, and many more through plugins[1]. Since the default format for QA4IE is GATE XML, we suggest using GATE as a preprocessing step to convert your annotations if they're in a different format. Here we use some synthetic clinical notes as an example to demonstrate the application of QA4IE to a sequence labeling task. This task is to extract mobility information from the clinical notes by three annotators. The mobility information is organized as shown in Figure 1.

Note mobility related annotations have a hierarchical structure, and that the schema file lists only the associated entities, sub-entities, and corresponding attributes, and does not specify the nested structure between entities and sub-entities. Therefore, we list all the entities and sub-entities in the configuration file with the identifier []. The dependency of the sub-entities and their parent entities is indicated by "sub_entities=". In Figure 2, the mobility entity contains three types of sub-entities: Action, Assistance, and Quantification. Based on the nested structure given in the configuration file, QA4IE can detect errors in the sub-entity that exceed its parent entity and output them to the quality assurance report. The types of entities/sub-entities that are allowed to overlap with each entity/sub-entity are listed after "overlaps=", and the types of entities/sub-entities that are not listed are not allowed to

overlap with the entity/sub-entity by default. In addition, we also specify the type of task in the configuration file, i.e., sequence labeling or text classification, and character encoding type.

Once the configuration file is ready, we can directly run QA4IE to perform the QA task. The main menu is shown in Figure 3. Regardless of the option selected, users have the flexibility to select specific text, annotators, entity/sub-entity, and features according to their needs.

Figure 2: Configuration file for mobility information extraction.

Figure 3: QA4IE main menu.

Document Validation includes two functions: "Text Differences" and "Set Name Difference". If errors have been found, users can generate reports to output the text and indices of text fragments where discrepancies have been found, as well as differences in set names.

Figure 4: Example of Document Validations

Annotation Validations have several sub-options. "Annotation Overlaps" detects overlapping entities/sub-entities that violate regulation. "Subentity Boundaries" finds sub-entities that are beyond the scope of their parent entities, while "Annotation Boundaries" checks if any annotation falls outside the annotation set. "Zero Length Annotations" and "Negative Length Annotations" catch entities/sub-entities whose end indices are less than or equal to the starting indices. "Document Scope" identifies annotations that are beyond the scope of the full text.

"Validate Schema" checks if there are annotations that are not defined in the schema.

Figure 5: Example of Annotation Validations

Statistics is used to collect information on various distributions, including the number of various types of entity/sub-entity and their lengths in terms of mean, variance, and extreme values. The choice of types can be either entity/sub-entity without regard to feature, or entity/sub-entity with a specific feature. Basic statistical information is displayed directly on the screen. A more complete version can be exported to the output folder

Figure 6: Example of Statistics

Evaluation provides both token-level and entity-level metrics because they provide complementary information.

Figure 7: Example of Evaluation

Figure 7 shows the token-level metric as an example, giving the confusion matrix, precision, recall, and F1 score for a selected type of entity/sub-entity.

Discrepancy Analysis is another highlight of QA4IE. Our Discrepancy Analysis has two unique features. As shown in Figure 8, QA4IE automatically categorizes discrepancies by text span and feature. This allows the annotation team to understand which are the main discrepancies that need to be resolved first.

Figure 8: Discrepancy Analysis

file	entity type	cluster	annotator1	annotator2	annotator3	discrepancy
11111-11111	Mobility	9	Mobility "Going down steps 4" (642-660) id=683 Subject : Patient Type : Objective Timeline : Present	Mobility "-Going down steps 4 -Squatting 5 -Sitting 15 min. 6" (641-693) id=684 Subject : Patient Type : Objective Timeline : Present	Mobility "Going down steps 4" (642-660) id=683 Subject : Patient Type : Objective Timeline : Present	length
11111-11111	Mobility	9	Mobility "Squatting 5" (663-674) id=685 Type : Objective Subject : Patient Timeline : Present		Mobility "Squatting 5" (663-674) id=685 Type : Objective Subject : Patient Timeline : Present	not annotated by all split combine split combine
11111-11111	Mobility	9	Mobility "Sitting 15 min. 6" (676-693) id=686 Type : Objective Subject : Patient Timeline : Present		Mobility "Sitting 15 min. 6" (676-693) id=686 Type : Objective Subject : Patient Timeline : Present	not annotated by all split combine

Table 1. Cluster based Discrepancy Analysis

In addition, unlike other annotation tools, QA4IE organizes discrepancies into clusters. As shown in Table 1, Annotator2 labels text with indices from 641 to 693 as one mobility entity, while Annotator1 and Annotator3 label them as three short mobility entities. Although the number of entities and the length of the text span are different, it is obvious that they should be compared together. This is the only way to correctly identify that the main discrepancy here is the split/combine issue. However, none of the existing tools can do this because they are all based on one-to-one correspondence. For example, the Annotation Difference tool in GATE can only compare the entities from 642 to 660 annotated by Annotator1 and Annotator3 with the long entity annotated by Annotator2, while treating the other two short entities as completely unrelated ones.

Due to space limitations, an exhaustive exposition of QA4IE's features and results are beyond the scope of this paper. More details, as well as the demonstration video, have been placed on GitHub at https://github.com/CC-RMD-EpiBio/QA4IE.

5. Related Work: Competitive Analysis

To the best of our knowledge, though all annotation tools include some semblance of QA, QA4IE is the first QA tool aimed to support QA of annotations for information extraction in a comprehensive manner.

We contextualize our contribution in light of a recent review paper that provides a comprehensive assessment and summary of the existing popular annotation tools (Neves and Ševa, 2021). The authors initially reviewed 784501

annotation tools and developed five principles, namely available, web-based, installable, workable, and schematic, to select annotation tools for detailed evaluation. Based on these five principles, 15 annotation tools were evaluated. Here we expand on that analysis by adding dimensions related to QA. Specifically, we include error checking, IRR, and discrepancy analysis. We then re-evaluate the same 15 tools according to our dimensions by referencing published papers, tool websites, and user guides. According to that analysis, not all of the tools included quality control facilities according to our definition, and only seven included at least one of the three core QA features. We also similarly evaluated six additional annotation tools using the same QA dimensions, and summarized the QA features of the thirteen annotation tools with QA features in Table 2, alongside QA4IE. To make it easier to understand the detailed features of the fourteen tools, including QA4IE, we have provided a GitHub page with the URLs of all these tools at https://github.com/CC-RMD-EpiBio/QA4IE.

Table 2 shows that most annotation tools do not have error checking capabilities. LightTag (LightTag, 2021), which is relatively rich in error checking features, only has some simple checks for formatting and consistency. QA4IE provides further checks for impermissible entity overlaps, annotation that violates nested structure, text inconsistency for the same document, and so on. It should be noted that Brat (Stenetorp et al., 2012) also provides a configuration file that is very feature-rich, but its configuration file is adopted to limit the choices of annotation to reduce errors, rather than checking for annotation errors.

Most annotation tools use IAA to evaluate the quality of annotations. For text classification, kappa coefficient is the standard metric. For sequence labeling, F1 score has been preferred. F1 score at the entity level treats all entities equally regardless of their length. However, the existing entity-level F1 score usually assigns an arbitrary score for partial matches, which makes the evaluation biased. This problem is particularly serious in tasks with long entities where partial matches occur more frequently. Therefore, QA4IE provides both token-level and entity-level F1 scores together. This can give a more comprehensive and fair evaluation.

As discussed earlier, QA4IE's visualization capabilities go beyond the discrepancy analysis of existing annotation tools, which typically offer only difference visualization based on one-to-one entity correspondence. When discrepancies are examined across multiple annotators, all several types of discrepancies are combined and therefore more challenging to characterize at this level. QA4IE provides a better cluster-based discrepancy analysis and automatic categorization.

6. Conclusion and future work

QA4IE is an accessible and feature-rich quality assurance tool for information extraction. It provides more customized error checking, especially for annotation structures and constraints. In terms of discrepancy analysis, QA4IE extends one-to-one correspondence-based difference visualization to cluster-based discrepancy analysis. The discrepancy categorization feature of QA4IE is very useful for annotation improvement prioritization.

Currently QA4IE only handles annotations in XML format, so we plan to extend it to other common annotation formats. We will also explore new IAA metrics for sequence labeling, as both token-level and entity-level IAA metrics are flawed and do not have chance agreement correction. Although the GUI interface of a tool is not typically considered a core feature of QA, we still plan to provide this option in the near future. We have shared the current QA4IE code on GitHub and will release new versions based on users' feedback.

7. Acknowledgements

Tool Name	Check & Validation							IAA		Discrepancy Analysis		
	Text diff	Set name	formatting					F1		1-1 diff	cluster diff	Discrepancy categorization
			Zero/neg length	overlap	scope	schema	boundary	Token	Entity			
GATE									X	X		
ERAS												
Brat									X			
LightTag		X	X						X	X		
MAT										X		
Tagtog								X				
TextAE								X				
Watson KS												
WebAnno			X									
Prodigy												
SLATE										X		
Knowtator								X	X			
INCEpTION												
QA4IE	X	X	X	X	X	X	X	X	X		X	X

Table 2. Comparison of the quality assurance function of QA4IE with other common annotation tools

8. Bibliographical References

Artstein, Ron, and Massimo Poesio. "Inter-coder agreement for computational linguistics." *Computational Linguistics* 34.4 (2008): 555-596.

Atdağ, Samet, and Vincent Labatut. "A comparison of named entity recognition tools applied to biographical texts." *2nd International conference on systems and computer science*. IEEE, 2013.

Azzopardi, Leif, Paul Thomas, and Alistair Moffat. "cwl_eval: An evaluation tool for information retrieval." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

Beasley, Lucas, and Prashanti Manda. *Comparison of natural language processing tools for automatic gene ontology annotation of scientific literature*. No. e27028v1. PeerJ Preprints, 2018.

Cejuela, Juan Miguel, et al. "tagtog: interactive and text-mining-assisted annotation of gene mentions in PLOS full-text articles." *Database* 2014 (2014).

Cunningham, Hamish, et al. "Getting more out of biomedical documents with GATE's full lifecycle open source text analytics." PLoS computational biology 9.2 (2013): e1002854.

Dai, Xiang. "Recognizing complex entity mentions: A review and future directions." *Proceedings of ACL 2018, Student Research Workshop*. 2018.

de Castilho, Richard Eckart, et al. "WebAnno: a flexible, web-based annotation tool for CLARIN." *Proceedings of the CLARIN Annual Conference (CAC)*. 2014.

Grosman, Jonatas S., et al. "Eras: Improving the quality control in the annotation process for natural language processing tasks." *Information Systems* 93 (2020): 101553.

Grundke, Maximilian, et al. "TextAI: Enhancing TextAE with Intelligent Annotation Support." *SMBM*. 2016.

Ide, Nancy, and James Pustejovsky, eds. *Handbook of linguistic annotation*. Springer, 2017.

Kim, Jin-Dong, et al. "Open Agile text mining for bioinformatics: the PubAnnotation ecosystem." *Bioinformatics* 35.21 (2019): 4372-4380.

Klie, Jan-Christoph, et al. "The inception platform: Machine-assisted and knowledge-oriented interactive annotation." *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*. 2018.

Kummerfeld, Jonathan K. "SLATE: a super-lightweight annotation tool for experts." *arXiv preprint arXiv:1907.08236* (2019).

Lenzi, Valentina Bartalesi, Giovanni Moretti, and Rachele Sprugnoli. "CAT: the CELCT Annotation Tool." *LREC*. 2012.

LightTag, "LightTag," Online, 2021. [Online]. Available: https://lighttag.io

Neves, Mariana, and Jurica Ševa. "An extensive review of tools for manual annotation of documents." *Briefings in bioinformatics* 22.1 (2021): 146-163.

Nadeau, David, and Satoshi Sekine. "A survey of named entity recognition and classification." *Lingvisticae Investigationes* 30.1 (2007): 3-26.

Ogren, Philip. "Knowtator: a protégé plug-in for annotated corpus construction." *Proceedings of the Human*

Language Technology Conference of the NAACL, Companion Volume: Demonstrations. 2006.

Pustejovsky, James, and Amber Stubbs. *Natural Language Annotation for Machine Learning: A guide to corpus-building for applications*. " O'Reilly Media, Inc.", 2012.

Shindo, Hiroyuki, Yohei Munesada, and Yuji Matsumoto. "PDFAnno: a web-based linguistic annotation tool for pdf documents." *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. 2018

Stenetorp, Pontus, et al. "BRAT: a web-based tool for NLP-assisted text annotation." *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*. 2012.