

YEDDA: A Lightweight Collaborative Text Span Annotation Tool

Jie Yang, Yue Zhang, Linwei Li, Xingxuan Li

Singapore University of Technology and Design

{jie_yang, linwei_li, xingxuan_li}@mymail.sutd.edu.sg

yue_zhang@sutd.edu.sg

Abstract

In this paper, we introduce YEDDA, a lightweight but efficient and comprehensive open-source tool for text span annotation. YEDDA provides a systematic solution for text span annotation, ranging from collaborative user annotation to administrator evaluation and analysis. It overcomes the low efficiency of traditional text annotation tools by annotating entities through both command line and shortcut keys, which are configurable with custom labels. YEDDA also gives intelligent recommendations by learning the up-to-date annotated text. An administrator client is developed to evaluate annotation quality of multiple annotators and generate detailed comparison report for each annotator pair. Experiments show that the proposed system can reduce the annotation time by half compared with existing annotation tools. And the annotation time can be further compressed by 16.47% through intelligent recommendation.

1 Introduction

Natural Language Processing (NLP) systems rely on large-scale training data (Marcus et al., 1993) for supervised training. However, manual annotation can be time-consuming and expensive. Despite detailed annotation standards and rules, inter-annotator disagreement is inevitable because of human mistakes, language phenomena which are not covered by the annotation rules and the ambiguity of language itself (Plank et al., 2014).

Existing annotation tools (Cunningham et al., 2002; Morton and LaCivita, 2003; Chen and Styler, 2013; Druskat et al., 2014) mainly focus on providing a visual interface for user annotation

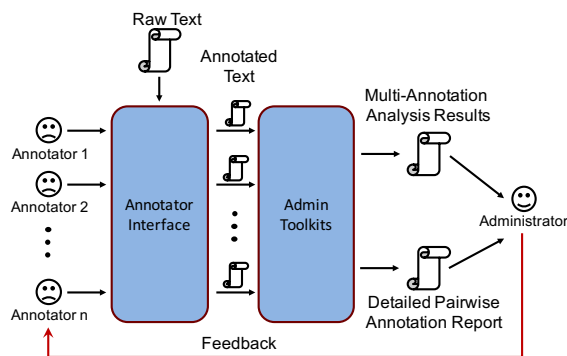


Figure 1: Framework of YEDDA.

process but rarely consider the post-annotation quality analysis, which is necessary due to the inter-annotator disagreement. In addition to the annotation quality, efficiency is also critical in large-scale annotation task, while being relatively less addressed in existing annotation tools (Ogren, 2006; Stenetorp et al., 2012). Besides, many tools (Ogren, 2006; Chen and Styler, 2013) require a complex system configuration on either local device or server, which is not friendly to new users.

To address the challenges above, we propose YEDDA¹, a lightweight and efficient annotation tool for text span annotation. A snapshot is shown in Figure 2. Here text span boundaries are selected and assigned with a label, which can be useful for Named Entity Recognition (NER) (Tjong Kim Sang and De Meulder, 2003), word segmentation (Sproat and Emerson, 2003), chunking (Tjong Kim Sang and Buchholz, 2000), etc. To keep annotation efficient and accurate, YEDDA provides systematic solutions across the whole annotation process, which includes the shortcut annotation, batch annotation with a command line, intelligent recommendation, format exporting and

¹Code is available at <https://github.com/jiesutd/YEDDA>.

Tool	Operating System			Self Consistency	Command Line Annotation	System Recommendation	Analysis	Size	Language
	MacOS	Linux	Win						
WordFreak	✓	✓	✓	✓	×	✓	×	1.1M	Java
GATE	✓	✓	✓	✓	×	✓	×	544M	Java
Knowtator	✓	✓	✓	×	×	×	✓	1.5M	Java
Stanford	✓	✓	✓	✓	×	×	×	88k	Java
Atomic	✓	✓	✓	✓	×	×	×	5.8M	Java
WebAnno	✓	✓	✓	×	×	✓	✓	13.2M	Java
Anafora	✓	✓	×	×	×	×	×	1.7M	Python
BRAT	✓	✓	×	×	×	✓	×	31.1M	Python
YEDDA	✓	✓	✓	✓	✓	✓	✓	80k	Python

Table 1: Annotation Tool Comparison .

administrator evaluation/analysis.

Figure 1 shows the general framework of YEDDA. It offers annotators with a simple and efficient Graphical User Interface (GUI) to annotate raw text. For the administrator, it provides two useful toolkits to evaluate multi-annotated text and generate detailed comparison report for annotator pair. YEDDA has the advantages of being:

- **Convenient:** it is lightweight with an intuitive interface and does not rely on specific operating systems or pre-installed packages.
- **Efficient:** it supports both shortcut and command line annotation models to accelerate the annotating process.
- **Intelligent:** it offers user with real-time system suggestions to avoid duplicated annotation.
- **Comprehensive:** it integrates useful toolkits to give the statistical index of analyzing multi-user annotation results and generate detailed content comparison for annotation pairs.

This paper is organized as follows: Section 2 gives an overview of previous text annotation tools and the comparison with ours. Section 3 describes the architecture of YEDDA and its detail functions. Section 4 shows the efficiency comparison results of different annotation tools. Finally, Section 5 concludes this paper and give the future plans.

2 Related Work

There exists a range of text span annotation tools which focus on different aspects of the annotation process. Stanford manual annotation tool² is a lightweight tool but does not support result analysis and system recommendation. Knowtator (Ogren, 2006) is a general-task annotation tool which links to a biomedical onto ontology to

²<http://nlp.stanford.edu/software/stanford-manual-annotation-tool-2004-05-16.tar.gz>

help identify named entities and relations. It supports quality control during the annotation process by integrating simple inter-annotator evaluation, while it cannot figure out the detailed disagreed labels. WordFreak (Morton and LaCivita, 2003) adds a system recommendation function and integrates active learning to rank the unannotated sentences based on the recommend confidence, while the post-annotation analysis is not supported.

Web-based annotation tools have been developed to build operating system independent annotation environments. GATE³ (Bontcheva et al., 2013) includes a web-based with collaborative annotation framework which allows users to work collaboratively by annotating online with shared text storage. BRAT (Stenetorp et al., 2012) is another web-based tool, which has been widely used in recent years, it provides powerful annotation functions and rich visualization ability, while it does not integrate the result analysis function. Anafora (Chen and Styler, 2013) and Atomic (Druskat et al., 2014) are also web-based and lightweight annotation tools, while they don't support the automatic annotation and quality analysis either. WebAnno (Yimam et al., 2013; de Castilho et al., 2016) supports both the automatic annotation suggestion and annotation quality monitoring such as inter-annotator agreement measurement, data curation, and progress monitoring. It compares the annotation disagreements only for each sentence and shows the comparison within the interface, while our system can generate a detailed disagreement report in .pdf file through the whole annotated content. Besides, those web-based annotation tools need to build a server through complex configurations and some of the servers cannot be deployed on Windows systems.

³GATE is a general NLP tool which includes annotation function.

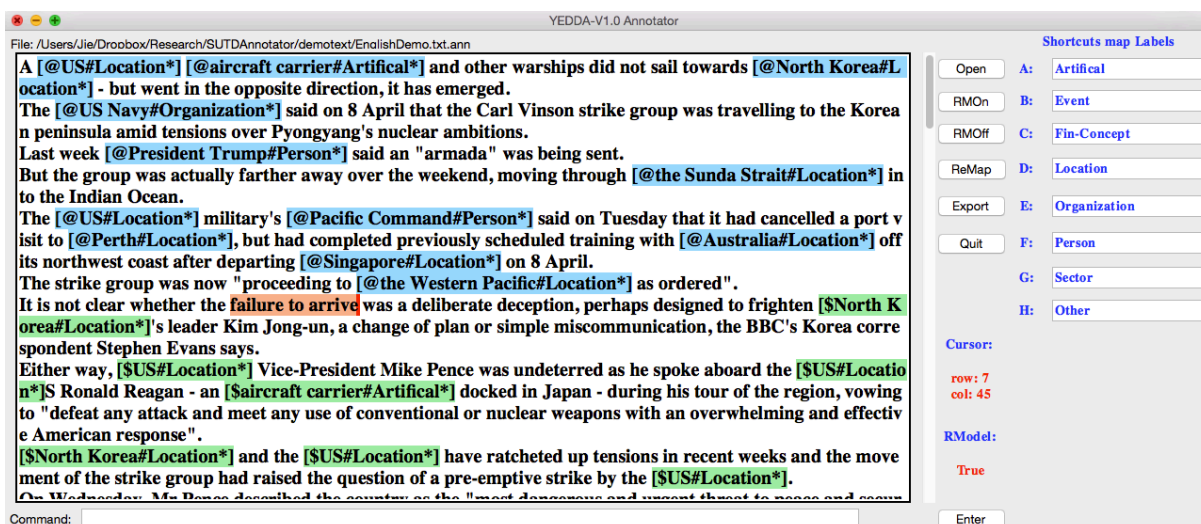


Figure 2: Annotator Interface.

The differences between YEDDA and related work are summarised in Table 1⁴. Here “Self Consistency” represents whether the tool works independently or it relies on pre-installed packages. Compared to these tools, YEDDA provides a lighter but more systematic choice with more flexibility, efficiency and less dependence on system environment for text span annotation. Besides, YEDDA offers administrator useful toolkits for evaluating the annotation quality and analyze the detailed disagreements within annotators.

3 YEDDA

YEDDA is developed based on standard Python GUI library Tkinter⁵, and hence needs only Python installation as a prerequisite and is compatible with all Operating System (OS) platforms with Python installation. It offers two user-friendly interfaces for annotators and administrator, respectively, which are introduced in detail in Section 3.1 and Section 3.2, respectively.

3.1 Annotator Client

The client is designed to accelerate the annotation process as much as possible. It supports shortcut annotation to reduce the user operation time. Command line annotation is designed to annotate multi-span in batch. In addition, the client provides system recommendations to lessen the workload of duplicated span annotation.

⁴For web-based tools, we list their server-side dependency on operating systems.

⁵<https://docs.python.org/2/library/tkinter.html>

Figure 2 shows the interface of annotator client on an English entity annotation file. The interface consists of 5 parts. The working area in the up-left which shows the texts with different colors (blue: annotated entities, green: recommended entities and orange: selected text span). The entry at the bottom is the command line which accepts annotation command. There are several control buttons in the middle of the interface, which are used to set annotation model. The status area is below the control buttons, it shows the cursor position and the status of recommending model. The right side shows the shortcut map, where shortcut key “a” or “A” means annotating the text span with “Artificial” type and the same for other shortcut keys. The shortcut map can be configured easily⁶. Details are introduced as follows.

3.1.1 Shortcut Key Annotation

YEDDA provides the function of annotating text span by selecting using mouse and press shortcut key to map the selection into a specific label. It is a common annotation process in many annotation tools (Stenetorp et al., 2012; Bontcheva et al., 2013). It binds each label with one custom shortcut key, this is shown in the “Shortcuts map Labels” part of Figure 2. The annotator needs only two steps to annotate one text span, i.e. “select and press”. The annotated file updates simultaneously with each key pressing process.

⁶Type the corresponding labels into the entries following shortcut keys and press “ReMap” button.

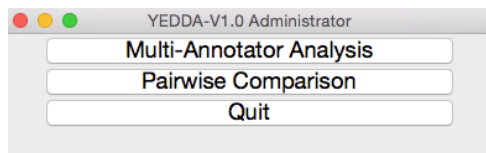


Figure 3: Administrator Interface.

3.1.2 Command Line Annotation

YEDDA also support the command line annotation function (see the command entry in the bottom of Figure 2) which can execute multi-span annotation at once. The system will parse the command automatically and convert the command into multi-span annotation instructions and execute in batch. It is quite efficient for the tasks of character-based languages (such as Chinese and Japanese) with high entity density. The command follows a simple rule which is “ $n1 + key1 + n2 + key2 + n3 + key3 + \dots$ ”, where ‘ $n1, n2, n3$ ’ are the length of the entities and ‘ $key1, key2, key3$ ’ is the corresponding shortcut key. For example, command “ $2A3D2B$ ” represents annotating following 2 characters as label ‘ A ’ (mapped into a specific label name), the following 3 characters as label ‘ D ’ and 2 characters further as label ‘ B ’.

3.1.3 System Recommendation

It has been shown that using pre-annotated text and manual correction increases the annotation efficiency in many annotation tasks (Meurs et al., 2011; Stenetorp et al., 2012). YEDDA offers annotators with system recommendation based on the existing annotation history. The current recommendation system incrementally collects annotated text spans from sentences that have been labeled, thus gaining a dynamically growing lexicon. Using the lexicon, the system automatically annotates sentences that are currently being annotated by leveraging the forward maximum matching algorithm. The automatically suggested text spans and their types are returned with colors in the user interface, as shown in green in Figure 2. Annotators can use the shortcut to confirm, correct or veto the suggestions. The recommending system keeps online updating during the whole annotation process, which learns the up-to-date and in-domain annotation information. The recommending system is designed as “pluggable” which ensures that the recommending algorithm can be easily extended into other sequence labeling mod-

F:Entity/Chunk	UserA	UserB	UserC	UserD	UserE
UserA	100/100	64.1/74.5	70.5/80.9	52.6/62.5	59.0/67.3
UserB	64.1/74.5	100/100	64.8/74.9	58.8/67.1	56.1/66.7
UserC	70.5/80.9	64.8/74.9	100/100	54.7/64.6	57.9/69.2
UserD	52.6/62.5	58.8/67.1	54.7/64.6	100/100	63.7/75.9
UserE	59.0/67.3	56.1/66.7	57.9/69.2	63.7/75.9	100/100

Figure 4: Multiple annotators F1-score matrix.

els, such as Conditional Random Field (CRF)⁷ (Lafferty et al., 2001). The recommendation can be controlled through two buttons “RMon” and “RMOff”, which enables and disables the recommending function, respectively.

3.1.4 Annotation Modification

It is inevitable that the annotator or the recommending system gives incorrect annotations or suggestions. Based on our annotation experience, we found that the time cost of annotation correction cannot be neglected. Therefore, YEDDA provides several efficient modification actions to revise the annotation:

- **Action withdraw:** annotators can cancel their previous action and let system return to the last status by press the shortcut key `Ctrl+z`.
- **Span label modification:** if the selected span has the correct boundary but receives an incorrect label, annotator only needs to put the cursor inside the span (or select the span) and press the shortcut key of the right label to correct label.
- **Label deletion:** similar to the label modification, the annotator can put the cursor inside the span and press shortcut key `q` to remove the annotated (recommended) label.

3.1.5 Export Annotated Text

As the annotated file is saved in `.ann` format, YEDDA provides the “Export” function which exports the annotated text as standard format (ended with `.anns`). Each line includes one word/character and its label, sentences are separated by an empty line. The exported label can be chosen in either BIO or BIOES format (Ratinov and Roth, 2009).

3.2 Administrator Toolkits

For the administrator, it is important and necessary to evaluate the quality of annotated files and analyze the detailed disagreements of different annotators. Shown in Figure 3, YEDDA provides a

⁷Those sequence labeling models work well on big training data. For limited training data, the maximum matching algorithm gives better performance.

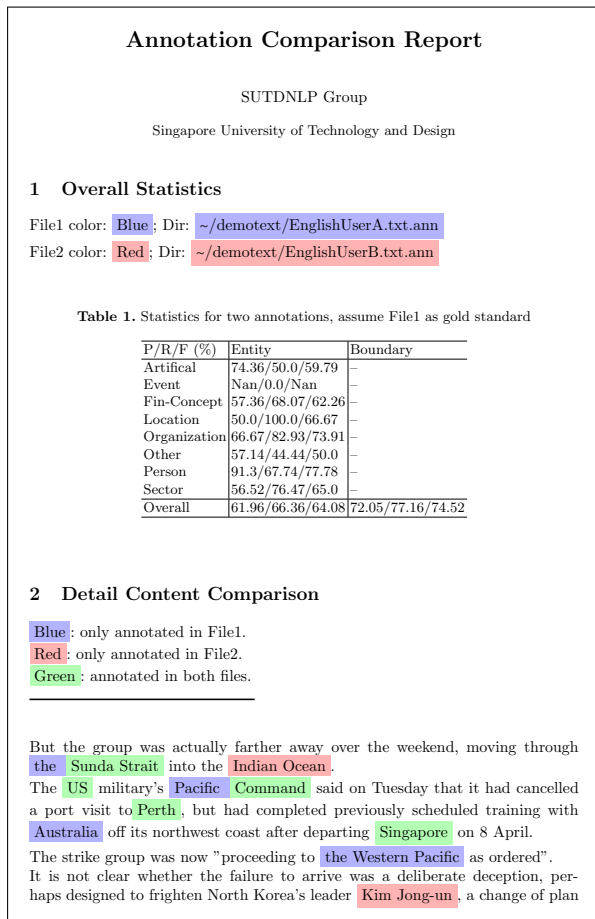


Figure 5: Detailed report for annotator pair.

simple interface with several toolkits for administrator monitoring the annotation process.

3.2.1 Multi-Annotator Analysis

To evaluate and monitor the annotation quality of different annotators, our Multi-Annotator Analysis (MAA) toolkit imports all the annotated files and gives the analysis results in a matrix. As shown in Figure 4, the matrix gives the F1-scores in full level (consider both boundary and label accuracy) and boundary level (ignore the label correctness, only care about the boundary accuracy) of all annotator pairs.

3.2.2 Pairwise Annotators Comparison

If an administrator wants to look into the detailed disagreement of annotators, it is quite convenient by using the Pairwise Annotators Comparison (PAC). PAC loads two annotated files and generates a specific comparison report file⁸ for the two annotators. As shown in Figure 5, the report

⁸The report is generated in `.tex` format and can be compiled into `.pdf` file.




Figure 6: Speed comparison.

is mainly in two parts:

- **Overall statistics:** it shows the specific precision, recall and F1-score⁹ of two files in all labels. It also gives the three accuracy indexes on overall full level and boundary level in the end.
- **Content comparison:** this function gives the detailed comparison of two annotated files in whole content. It highlights the annotated parts of two annotators and assigns different colors for the agreed and disagreed span.

4 Experiments

Here we compare the efficiency of our system with four widely used annotation tools. We extract 100 sentences from CoNLL 2003 English NER (Tjong Kim Sang and De Meulder, 2003) training data, with each sentence containing at least 4 entities. Two undergraduate students without any experience on those tools are invited to annotate those sentences¹⁰. Their average annotation time is shown in Figure 6, where "YEDDA+R" suggests annotation using YEDDA with the help of system recommendation. The inter-annotator agreements for those tools are closed, which around 96.1% F1-score. As we can see from the figure, our YEDDA system can greatly reduce the annotation time. With the help of system recommendation, the annotation time can be further reduced. We notice that "YEDDA+R" has larger advantage with the increasing numbers of annotated sentences, this is because the system recommendation gives better

⁹Notice that we assume "File1" as a gold standard. This only affects the order of precision and recall, while the F1-score keeps same if we choose the other file as gold standard.

¹⁰We ask the students to annotate those sentences several rounds to get familiar with the entities before they start the final exercise with recording.

suggestions when it learns larger annotated sentences. The “YEDDA+R” gives 16.47% time reduction in annotating 100 sentences¹¹.

5 Conclusion and Future Work

We have presented a lightweight but systematic annotation tool, YEDDA, for annotating the entities in text and analyzing the annotation results efficiently. In order to reduce the workload of annotators, we are going to integrate active learning strategy in our system recommendation part in the future. A supervised sequence labeling model (such as CRF) is trained based on the annotated text, then unannotated sentences with less confidence (predicted by this model) are reordered in the front to ensure annotators only annotate the most confusing sentences.

6 Acknowledgements

We thank Yanxia Qin, Hongmin Wang, Shaolei Wang, Jiangming Liu, Yuze Gao, Ye Yuan, Lu Cao, Yumin Zhou and other members of SUTDNLP group for their trials and feedbacks. Yue Zhang is the corresponding author. Jie is supported by the YEDDA grant 52YD1314.

References

- Kalina Bontcheva, Hamish Cunningham, Ian Roberts, Angus Roberts, Valentin Tablan, Niraj Aswani, and Genevieve Gorrell. 2013. Gate teamware: a web-based, collaborative text annotation framework. *Language Resources and Evaluation* 47(4):1007–1029.
- Wei-Te Chen and Will Styler. 2013. Anafora: a web-based general purpose annotation tool. In *NAACL* volume 2013, page 14.
- Hamish Cunningham, Diana Maynard, Kalina Bontcheva, and Valentin Tablan. 2002. Gate: an architecture for development of robust hlt applications. In *ACL*, pages 168–175.
- Richard Eckart de Castilho, Eva Mujdricza-Maydt, Seid Muhie Yimam, Silvana Hartmann, Iryna Gurevych, Anette Frank, and Chris Biemann. 2016. A web-based tool for the integrated annotation of semantic and syntactic structures. In *Proceedings of the Workshop on LT4DH*, pages 76–84.
- Stephan Druskat, Lennart Bierkandt, Volker Gast, Christoph Rzymiski, and Florian Zipser. 2014. Atomic: An open-source software platform for multi-level corpus annotation. In *Proceedings of the 12th Konferenz zur Verarbeitung natürlicher Sprache*.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, volume 951, pages 282–289.
- Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Computational linguistics* 19(2):313–330.
- Marie-Jean Meurs, Caitlin Murphy, Nona Naderi, Ingo Morgenstern, Carolina Cantu, Shary Semarjit, Greg Butler, Justin Powlowski, Adrian Tsang, and René Witte. 2011. Towards evaluating the impact of semantic support for curating the fungus scientific literature. *WS2* page 34.
- Thomas Morton and Jeremy LaCivita. 2003. Wordfreak: an open tool for linguistic annotation. In *NAACL: Demo*, pages 17–18.
- Philip V Ogren. 2006. Knowtator: a protégé plugin for annotated corpus construction. In *NAACL: Demo*, Association for Computational Linguistics, pages 273–275.
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. Linguistically debatable or just plain wrong? In *ACL*, volume 2, pages 507–511.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *CoNLL*, pages 147–155.
- Richard Sproat and Thomas Emerson. 2003. The first international chinese word segmentation bake-off. In *Proceedings of the SIGHAN workshop on CLP*, pages 133–143.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun’ichi Tsujii. 2012. Brat: a web-based tool for nlp-assisted text annotation. In *EACL: Demo*, pages 102–107.
- Erik F Tjong Kim Sang and Sabine Buchholz. 2000. Introduction to the conll-2000 shared task: Chunking. In *Proceedings of the 2nd workshop on Learning language in logic and the 4th conference on Computational natural language learning-Volume 7*, pages 127–132.
- Erik F Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *HLT-NAACL*, Association for Computational Linguistics, pages 142–147.
- Seid Muhie Yimam, Iryna Gurevych, Richard Eckart de Castilho, and Chris Biemann. 2013. Webanno: A flexible, web-based and visually supported system for distributed annotations. In *ACL: Demo*, pages 1–6.

¹¹The speed improvement by recommendation depends on the density of text spans. We suggest enabling the recommendation model in the task whose text contains dense and recurring text spans.