

WritingAssistant™

Learn at your pace, in your context

Revised: March 2012

Introduction

English is well known as the business language of the world – it is estimated that about one-fourth of the world’s population can communicate, at least to a small degree, in English. Its willingness to accept words and even grammatical structures from the vernacular has facilitated its rapid spread as the world’s lingua franca. While enabling people separated by great distances and cultures to communicate easily, it also makes business more challenging for those who have not yet mastered the language and creates a hierarchy, ranking those with greater familiarity over those with less.

People who learn English in a classroom but do not hear it spoken in the context in which they will have to use the language cannot reinforce the skills they have learned in class until they are thrown, often unprepared, into a situation that demands it. There is need for a tool that will detect issues in ESL writing comprehensively - one that is non-intrusive, can be contextually relevant, and is not static.

The WritingAssistant enables writers of all levels to improve their writing skills in a non-intrusive, contextually relevant manner by allowing users of all levels to check any writing sample for grammar, style and content quality. Currently on a web portal, it provides easy access to a proofreading and self-learning tool that can be used anywhere and in any context.

The WA Taxonomy of Errors

The WritingAssistant checks text for errors in three broad categories: Mechanics, Grammar, and Style. The mechanics category includes spelling, punctuation and case, looking for words that are misspelled or incorrectly capitalized, as well as checking for clear and correct usage of punctuation. The grammar rules find errors in syntax and tense, as well as pointing out words that might be inappropriately chosen. Style covers the overall flow of the sentences, checking if they adhere to conventions of good writing.

Mechanics	Grammar	Vocabulary	Content Quality
<ul style="list-style-type: none">• Punctuation• Spelling• Capitalization• Spacing	<ul style="list-style-type: none">• Subject-Verb Agreement• Other Verb Problems• Incorrect Preposition• Article Usage• Singular/Plural Issues• Sentence Fragments• Run-on Sentences• Pronoun-Antecedent Agreement	<ul style="list-style-type: none">• SMS Lingo• Indian English• Incorrect Word Usage• Wordy Sentences	<ul style="list-style-type: none">• Relevance• Word Usage Level

The numbers and types of errors are tallied to create an assessment score of the text. The subcategories can be made optional depending on user requirements; if the user desires that only grammar and not style should be examined, or even finer subcategories such as Article Usage within grammar should be omitted, this can be included in the user package. The weight of each subcategory in the final score can be customized so that an overall assessment rating reflects features that the user wishes to emphasize.

Architecture

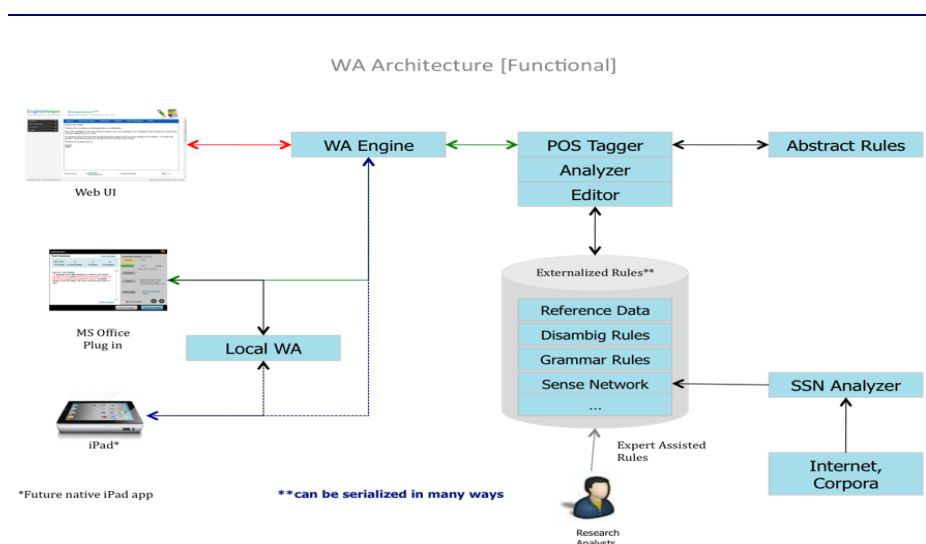
WA is designed to be highly scalable and flexible. It benefits from a unique model-driven architecture that allows it to learn continuously. Experts can enhance many aspects of the WA Engine without needing to resort to programming. The engine’s architecture also allows it to be configured for different contexts rather easily.

The WA Engine comprises four logical components: a POS Tagger, an Analyzer, an Editor and the SSN Analyzer. The POS Tagger determines the POS of the words in the text. The Analyzer analyses the text and the Editor determines the corrections for any errors that are found. The SSN Analyzer analyzes text and enhances the semantic sense network for the corresponding vocabulary.

The WA engine also allows for the assessment of content quality in two ways: the first is by allowing the writer to work within a predetermined contextual framework, and the second is by determining context through a semantic sense network (SSN).

POS tagging, assessment of grammar and determining relevance are all highly dependent on understanding the context of a written piece of work. The WA has functionality that allows a 'guided form' in which an administrator can set a writing prompt. The administrator can then also add multiple 'gold standard responses' to the project, which function as ideal responses to the given prompt. The engine then uses the responses and machine learning techniques to generate a list of key concepts in the project. This list can assess the relevance of a writing sample submitted by a user to the prompt. The words in the list can also be added to the custom dictionary, allowing them to take on a different meaning or new properties from the dictionary meaning of the terms, e.g., to denote proper nouns, definiteness for article usage, etc..

WA can be deployed in several form factors – on the web, as an MS Office plug in and in the future will be available on tablet devices in a native form. When deployed as a plug-in, a synchronization process ensures that the rules in the local version are updated with any changes on the WA Server.



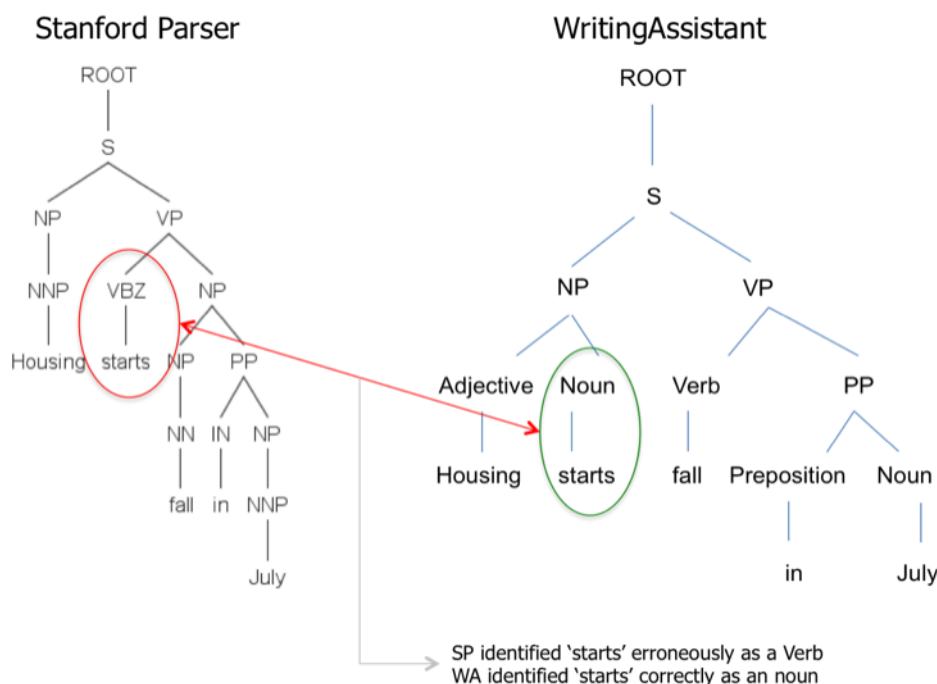
POS Tagger

The Part of Speech (POS) Tagger is the first major step to parsing and correcting the submitted text. Each word is assigned a part of speech based on its function in the sentence (noun, verb, etc.). The POS Tagger has been created using corpora-independent methods, so that there is no rule bias created by the confines of a text corpus.

The rules of the tagger encapsulate the patterns of normal English syntax, but are flexible enough to allow for variation from imperfect user input. They determine the word's POS based on specific context that speakers would use to disambiguate the word's function. The POS-tagger also uses N-grams, which are common constructions that follow a fixed pattern and might have a different function as a group than their individual components, possibly affecting the parsing of the sentence.

Because the Tagger is comprised of sets of externalized rules, and the SSN, it allows us to achieve higher accuracy in tagging across heterogeneous user bases than a probabilistic determination would. The rules allow the POS Tagger to be sensitive to recurring error patterns in the writing of ESL learners. The WA POS Tagger's accuracy competes with well-regarded taggers on tagging accuracy of the Brown corpus, and outperforms other taggers on learner data since those have been derived from corpora of well-formed text.

An example of the error produced most likely because of corpora dependency can be seen in the parse trees below for the sentence:
"Housing starts fall in July"



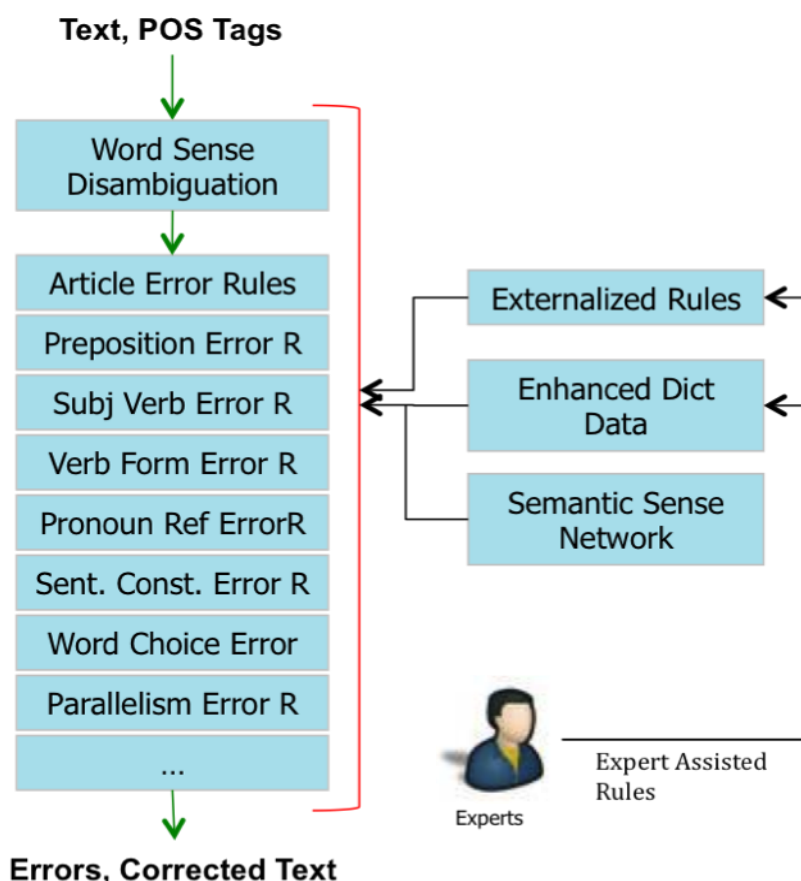
Analyzer and Editor

The Analyzer and Editor perform the type of analysis a human editor would upon reading the text. The Analyzer scans for grammatical errors in the text, focusing on the types of errors encountered in the writing of ESL learners. The Analyzer and Editor, like the POS Tagger, have a structure that allows for great flexibility: there is an underlying grammar of correct English on top of which we have a set of built-in abstract rules that detect deviations from the structure as well as an externalized rule dataset that looks for certain error patterns. These external rules can be customized for ESL learners coming from various base languages.

We have so far aimed at recognizing errors committed by Indian English learners. For example, since Indian languages have nothing similar to the articles "a, an, and the," they often struggle with using them appropriately and may leave them out altogether. This differs from ESL learners whose native language is Spanish or German, who understand the concept of 'definiteness' that is conveyed by using the various article forms and may have to learn only the idiomatic uses.

The Analyzer searches for these errors while the Editor attempts to correct the errors. Sometimes an obvious correction suggests itself, as in the case of Article Usage, when one simply has to use the other article. In other, more complicated cases such as Mixed Construction, the Analyzer may detect that the grammatical structure is inappropriate, but there is no obvious single way to correct the error. There are three main categories of errors – Mechanics, Grammar and Style - and subcategories within each that detect errors, especially those that can alter meaning and hamper communication.

Analyzer and Editor

**Effectiveness: Evidence from real world data**

We have been assessing WA's comparative performance relative to commercial automated grammar detection software using both well formed text and learner data acquired from India. It is clear that WA performs significantly better than any of these tools. Because the WA is much more comprehensive in error detection and correction than other grammar checkers of its kind, is targeted towards a general population of learners rather than as a remedial tool, and is fundamentally designed for non-native writers as opposed to native writers, the WritingAssistant is very effective for ESL learners.

We believe that WA's flexible architecture allows the product to continually improve as we observe its performance in real world settings. Below, we present the results of a couple of comparative tests.

Writing Assignment

The following writing assignment was assigned to a group of Indian English language learners:

"You are an employee for Bluewater Internet Bookstore. Write an apology email to a customer, John McNeal, who ordered a book, "Who Moved My Cheese," two weeks ago and has not received it yet."

A sample writing response to the prompt was:

Dear Mr.John,

I being a one of the DIRECTOR and a responsible person in blue water,apologize you personally for the delay.This was due to the shipping problem which we were unaware,And our ships could not reach the location at scheduled time as the weather condition dint support for sailing.hope you bear with us,am sure this wont happen again in future.

your's sincerely,

DIRECTOR(Blue water)

Comparative Results:

Corrections:	
WritingAssistant:	<p>Dear John, I being one of the directors and a responsible person at Bluewater apologize ?? you personally for the delay. This was due to the shipping problem of which we were unaware, and our ships could not reach the location at the scheduled time as the weather condition did not support sailing. Hope you bear with us, I am sure this will not happen again in future.</p> <p>Yours sincerely, DIRECTOR (Bluewater)</p>
Ginger:	<p>Dear Mr. John, I am being a one of the DIRECTOR and a responsible person in blue water, apologize you personally for the delay. This was due to the shipping problem which we were unaware, And our ships could not reach the location at the scheduled time as the weather condition dint support for sailing. Hope you bear with us, am sure this won't happen again in future.</p> <p>Your's sincerely, DIRECTOR (Blue water)</p>
Grammarly:	<p>Dear Mr. John, I being a one of the DIRECTOR and a responsible person in blue water, apologize you personally for the delay. This was due to the shipping problem which we were unaware. And our ships could not reach the location at scheduled time as the weather condition dint support for sailing. hope you bear with us, am sure this wont happen again in future.</p> <p>your sincerely, DIRECTOR(Blue water)</p>
Microsoft Word 2010:	<p>Dear Mr. John, I being a one of the DIRECTOR and a responsible person in blue water, apologize ?? you personally for the delay. This was due to the shipping problem, which we were unaware, and our ships could not reach the location at scheduled time as the weather condition dint support for sailing. Hope you bear with us, am sure this will not happen again in future.</p> <p>Yours sincerely, DIRECTOR (Blue water)</p>

Out of 27 errors, the WA found the most with an 88% accuracy rate on this example. The rest found fewer than 50% of the errors in the text.

Categories	Gold Standard	WA	Ginger	Grammarly	MS Word 2010
Article Usage	2	2	1		
Incorrect Case	5	4	1		
Incorrect Preposition	2	2			
Mixed Construction	1	1			
Other Verb Problems	2	2			
Punctuation	8	8	7	5	7
Run-on Sentences	1				
Singular/Plural Issues	2	1			
Spelling	4	4	2	2	3
Grand Total	27	24	11	7	10

Analysis of learner data

The following results were determined by examining 100 emails written by Indian English language learners. Three independent assessments of the errors found in the files were conducted and reconciled to create a master list of errors [gold standard]. A custom dictionary for each of the projects was created in the WritingAssistant. The results from each program on running these files were checked against the master list and are specified below:

Categories	Gold Standard	WA	Ginger	Grammarly	MS Word 2010
Content Quality	102	53%	12%	18%	5%
Mixed Construction	39	41%	26%	13%	13%
Wordy Sentences	63	60%	3%	21%	0%
Grammar	372	53%	20%	15%	7%
Article Usage	104	47%	25%	13%	2%
Incorrect Preposition	27	52%	7%	4%	0%
Indian English	27	81%	4%	4%	4%
Needed Words Missing	2	100%	0%	0%	0%
Other Verb Problems	84	70%	17%	20%	7%
Pronoun-Antecedent	6	50%	0%	0%	0%
Run-on Sentences	18	17%	0%	22%	11%
Sentence Fragment	18	67%	0%	0%	17%
Singular/Plural Issues	15	40%	47%	7%	27%
SMS Lingo	12	100%	100%	75%	17%
Subject-Verb Agreement	12	75%	33%	33%	33%
Verb Tense	47	15%	19%	11%	2%
Mechanics	965	91%	55%	41%	56%
Incorrect Case	317	91%	37%	10%	34%
Punctuation	402	89%	52%	47%	58%
Spelling	246	93%	84%	72%	83%
Grand Total	1,439	78%	43%	33%	40%

Summary of Results

The WritingAssistant found nearly twice as many errors overall as the other grammar checkers, with a much higher percentage of errors caught in those areas that English language learners struggle with. It also catches a much higher percentage of errors in areas such as punctuation and incorrect case because it can be customized to a particular topic by using the custom dictionary and by using the custom features to define the writing format. The WritingAssistant does not currently catch the same percentage of errors in all formats, but the percentages should improve as it is exposed to a larger and more varied sample set.

Conclusion

The WritingAssistant is a complete tool for analyzing and assessing writing. Its architecture allows for robust detection of errors, and is highly adaptable to different contexts and dialects. Because of this adaptability and because we have focused on the Indian English dialect, the WritingAssistant is also stronger than currently available online grammar checkers. Furthermore, the engine performs particularly well in the areas in which Indian English language learners struggle most, such as article usage, verb transitivity, and incorrect preposition usage. While the engine may not detect as many errors overall as human graders, it does avoid the problem of inter-rater reliability – it will detect the same number of errors in a specific example when run multiple times and the same types of errors across samples. The engine will also continue improving as we continue learning from new data submitted. We believe that the WritingAssistant solves the pressing need for a tool that assesses writing samples for mechanics, grammar, style and content quality, especially in the Indian context.