

Event-Centered Simplification of News Stories

Goran Glavaš

Faculty of Electrical
Engineering and Computing
University of Zagreb, Croatia
goran.glavas@fer.hr

Sanja Štajner

Research Group in
Computational Linguistics
University of Wolverhampton, UK
sanjastajner@wlv.ac.uk

Abstract

Newswire text is often linguistically complex and stylistically decorated, hence very difficult to comprehend for people with reading disabilities. Acknowledging that events represent the most important information in news, we propose an event-centered approach to news simplification. Our method relies on robust extraction of factual events and elimination of surplus information which is not part of event mentions. Experimental results obtained by combining automated readability measures with human evaluation of correctness justify the proposed event-centered approach to text simplification.

1 Introduction

For non-native speakers, people with low literacy or intellectual disabilities, and language-impaired people (e.g., autistic, aphasic, congenitally deaf) newswire texts are difficult to comprehend (Carroll et al., 1999; Devlin, 1999; Feng, 2009; Štajner et al., 2012). Making news equally accessible to people with reading disabilities helps their integration into society (Freyhoff et al., 1998).

In news, syntactically complex and stylistically decorated sentences, combining several pieces of information of varying relevance, are frequent. For example, in “*Philippines and China diplomatically resolved a tense naval standoff, the most dangerous confrontation between the sides in recent years.*” the “resolving of a standoff” is arguably a more relevant piece of information than the “standoff” being “the most dangerous confrontation in years”. However, studies indicate that people with reading disabilities, especially people with intellectual disabilities, have difficulties discriminating relevant

from irrelevant information (Pimperton and Nation, 2010), e.g., when sentences are particularly long and complex (Carretti et al., 2010; Feng, 2009). Thus, complex sentences need to be shortened and simplified, and any irrelevant content eliminated in order to reduce complexity of news stories and facilitate their comprehension.

News describes real-world events, i.e., events are dominant information concepts in news (Van Dijk, 1985; Pan and Kosicki, 1993). Although news is made up of event-oriented texts, the number of descriptive sentences and sentence parts relating to non-essential information is substantial (e.g., “*The South China Sea is home to a myriad of competing territorial claims*”). Such descriptions do not relate to any of the concrete events, but significantly contribute to the overall complexity of news.

Most existing approaches to text simplification address only lexical and syntactic complexity, i.e., they do not apply any content reduction (Carroll et al., 1998; Devlin and Unthank, 2006; Aluísio et al., 2008; Saggion et al., 2011). In this work we present a semantically-motivated, event-based simplification approach. We build upon state-of-the-art event extraction and discard text not belonging to extracted event mentions. We propose two event-based simplification schemes, allowing for different degrees of simplification. We evaluate event-centered simplification by combining automated measures of readability with human assessment of grammaticality and information relevance. Experimental results suggest that event-centered simplification is justified as it outperforms the syntactically-motivated baseline.

2 Related Work

Several projects dealt with automated text simplification for people with different reading difficulties:

people with alexia (Carroll et al., 1998; Devlin and Unthank, 2006), cognitive disabilities (Saggion et al., 2011), autism (Orasan et al., 2013), congenital deafness (Inui et al., 2003), and low literacy (Aluísio et al., 2008). Most of these approaches rely on rule-based lexical and syntactic simplification. Syntactic simplification is usually carried out by recursively applying a set of hand-crafted rules at a sentence level, not considering interactions across sentence boundaries. Lexical simplification usually substitutes difficult words with their simpler synonyms (Carroll et al., 1998; Lal and Ruger, 2002; Burstein et al., 2007).

Existing approaches dominantly rely on lexical and syntactic simplification, performing little content reduction, the exception being deletion of parenthetical expressions (Drndarevic et al., 2013). On the one hand, lack of content reduction has been recognized as one of the main shortcomings of automated systems (Drndarevic et al., 2013) which produce much worse simplification results compared to human. On the other hand, information extraction techniques help identify relevant content (e.g., named entities, events), but have not yet proven useful for text simplification. However, significant advances in event extraction (Ahn, 2006; Bethard, 2008; Llorens et al., 2010; Grover et al., 2010), achieved as the result of standardization efforts (Pustejovsky et al., 2003a; Pustejovsky et al., 2003b) and dedicated tasks (ACE, 2005; Verhagen et al., 2010), encourage event-oriented simplification attempts. To the best of our knowledge, the only reported work exploiting events for text simplification is that of Barlacchi and Tonelli (2013). They extract factual events from a set of Italian children’s stories and eliminate non-mandatory event arguments. They evaluate simplified texts using only the automated score which can hardly account for grammaticality and information relevance of the output.

We follow the idea of exploiting factual events for text simplification, acknowledging, however, that newswire texts are significantly more complex than children’s stories. Moreover, we complement automated readability measures with human assessment of grammaticality and information relevance. Furthermore, given that simplification systems often need to be tailored to the specific needs of a particular group (Orasan et al., 2013), and that people with different low literacy degrees need different levels of simplification (Scarton et al., 2010),

we offer two different simplification schemes. To the best of our knowledge, this is the first work on event-based text simplification for English.

3 Event-Centered Simplification

The simplification schemes we propose exploit the structure of extracted event mentions. We employ robust event extraction that involves supervised extraction of factual event anchors (i.e., words that convey the core meaning of the event) and the rule-based extraction of event arguments of coarse semantic types. Although a thorough description of the event extraction system is outside the scope of this paper, we describe the aspects relevant to the proposed simplification schemes.

3.1 Event Extraction

Our event extraction system performs supervised extraction of event anchors and a rule-based extraction of event arguments.

Anchor extraction. We use two supervised models, one for identification of event anchors and the other for classification of event type. The first model identifies tokens being anchors of event mentions (e.g., “resolved” and “standoff” in “*Philippines and China resolved a tense naval standoff.*”). The second model determines the TimeML event type (Pustejovsky et al., 2003a) for previously identified anchors. The models were trained with logistic regression using the following sets of features:

(1) *Lexical and PoS features* – word, lemma, stem, and PoS tag of the current token and the surrounding tokens (symmetric window of size 2);

(2) *Syntactic features* – the set of dependency relations and the chunk type (e.g., NP) of the current token. Additionally, we use features indicating whether the token governs nominal subject or direct object dependencies.

(3) *Modifier features* – modal modifiers (e.g., *might*), auxiliary verbs (e.g., *been*) and negations of the current token. These features help discriminate factual from non-factual events.

The supervised models were trained on the train portion of the EvExtra corpus¹, and tested on the separate test portion. The anchor identification model achieves precision of 83%, recall of 77%, and F-score performance of 80%. The model for event-type classification performs best for *Reporting* events, recognizing them with the F-score performance of 86%.

¹<http://takelab.fer.hr/data/grapheve/>

Table 1: Some of the patterns for argument extraction

Name	Example	Dependency relations	Arg. type
Nominal subject	“ China <u>confronted</u> Philippines”	<i>nsubj(confronted, China)</i>	Agent
Direct object	“China <u>disputes</u> the agreement ”	<i>doobj(disputes, agreement)</i>	Target
Prepositional object	“Philippines <u>protested</u> on Saturday ”; “The <u>confrontation</u> in South China Sea ”; “The <u>protest</u> against China ”	<i>prep(protested, on)</i> and <i>pobj(on, Saturday)</i> ; <i>prep(confrontation, in)</i> and <i>pobj(in, Sea)</i> ; <i>prep(protest, against)</i> and <i>pobj(against, China)</i>	Time Location Target
Participial modifier	“The vessel <u>carrying</u> missiles”; “The militant <u>killed</u> in the attack”	<i>partmod(vessel, carrying)</i> ; <i>partmod(militant, killed)</i>	Agent Target
Noun compound	“ Beijing <u>summit</u> ”; “ Monday <u>demonstrations</u> ”; “ UN <u>actions</u> ”	<i>nn(summit, Beijing)</i> ; <i>nn(demonstrations, Monday)</i> ; <i>nn(actions, UN)</i>	Location Time Agent

Argument extraction. We implement a rule-based extraction of event arguments, using a rich set of unlexicalized syntactic patterns on dependency parses as proposed in (Glavaš and Šnajder, 2013). All extraction patterns are defined with respect to event anchor and identify head words of arguments. We focus on extracting arguments of four coarse-grained types – *agent*, *target*, *time*, and *location* – for which we believe are informationally most relevant for the event. In total, there are 13 different extraction patterns, and their representative subset is presented in Table 1 (in examples, the argument is shown in bold and the anchor is underlined).

Some extraction patterns perform argument detection and classification simultaneously (e.g., a *nominal subject* is always an *agent*). Other patterns identify argument candidates, but further semantic processing is required to determine the argument type (e.g., *prepositional objects* can be *temporals*, *locations*, or *targets*). To disambiguate the argument type in such cases, we use named entity recognition (Finkel et al., 2005), temporal expression extraction (Chang and Manning, 2012), and WordNet-based semantic similarity (Wu and Palmer, 1994). Patterns based on dependency parse identify only the argument head words. The chunk of the argument head word is considered to be the full argument extent.

The argument extraction performance, evaluated on a held-out set, is as follows (F-score): agent – 88.0%, target – 83.1%, time – 82.3%, location – 67.5%.

3.2 Simplification Schemes

We base our simplification schemes on extracted event mentions. The rationale is that the most relevant information in news is made up of factual events. Thus, omitting parts of text that are not events would (1) reduce text complexity by eliminating irrelevant information and (2) increase readability by shortening long sentences. We propose two different simplification schemes:

(1) *Sentence-wise simplification* eliminates all the tokens of the original sentence that do not belong to any of the extracted factual event mentions (event anchors or arguments). A single sentence of the original text maps to a single sentence of the simplified text, assuming that the original sentence contains at least one factual event mention. Sentences that do not contain any factual event mentions (e.g., “*What a shame!*”) are removed from the simplified text. Algorithm 1 summarizes the sentence-wise simplification scheme.

(2) *Event-wise simplification* transforms each factual event mention into a separate sentence of the output. Since a single phrase can be an argument of multiple event mentions, a single input token may constitute several output sentences (e.g., “*China sent in its fleet and provoked Philippines*” is transformed into “*China sent in its fleet. China provoked Philippines.*”). We make three additional adjustments to retain the grammaticality of the output. Firstly, we ignore events of the *Reporting* type (e.g. *said*) as they frequently cannot constitute grammatically correct sentences on their own (e.g., “*Obama said.*”). Secondly, we do not

Algorithm 1. Sentence-wise simplification

```
input: sentence  $s$ 
input: set of event mentions  $\mathcal{E}$ 
// simplified sentence (list of tokens)
 $\mathcal{S} = \{\}$ 
// list of original sentence tokens
 $\mathcal{T} = \text{tokenize}(s)$ 
foreach token  $t$  in  $\mathcal{T}$  do
  foreach event mention  $e$  in  $\mathcal{E}$  do
    // set of event tokens
     $\mathcal{A} = \text{anchorAndArgumentTokens}(e)$ 
    // if the sentence token belongs to event
    if  $t$  in  $\mathcal{A}$  do
      // include token in simplified sentence
       $\mathcal{S} = \mathcal{S} \cup t$ 
    break
output:  $\mathcal{S}$ 
```

Algorithm 2. Event-wise simplification

```
input: sentence  $s$ 
input: set of event mentions  $\mathcal{E}$ 
// set of event-output sentence pairs
 $\mathcal{S} = \{\}$ 
// initialize output token set for each event
foreach  $e$  in  $\mathcal{E}$  do
   $\mathcal{S} = \mathcal{S} \cup (e, \{\})$ 
// list of original sentence tokens
 $\mathcal{T} = \text{tokenize}(s)$ 
foreach token  $t$  in  $\mathcal{T}$  do
  foreach event mention  $e$  in  $\mathcal{E}$  do
    // set of event tokens
     $a = \text{anchor}(e)$ 
     $\mathcal{A} = \text{anchorAndArgumentTokens}(e)$ 
    // part of verbal, non-reporting event
    if  $t$  in  $\mathcal{A}$  &  $\text{PoS}(a) \neq N$  &  $\text{type}(t) \neq \text{Rep}$  do
      // token is gerundive anchor
      if  $t = a$  &  $\text{gerund}(a)$ 
         $\mathcal{S}[e] = \mathcal{S}[e] \cup \text{pastSimple}(a)$ 
      else  $\mathcal{S}[e] = \mathcal{S}[e] \cup t$ 
output:  $\mathcal{S}$ 
```

transform events with nominal anchors into separate sentences, as such events tend to have very few arguments and are often arguments of verbal events. For example, in “*China and Philippines resolved a naval standoff*” mention “*standoff*” is a *target* of the mention “*resolved*”. Thirdly, we convert gerundive events that govern the clausal complement of the main sentence event into past simple for preserving grammaticality of the output. E.g., “*Philippines disputed China’s territorial claims, triggering the naval confrontation*” is transformed into “*Philippines disputed China’s territorial claims. Philippines triggered the naval confrontation.*”, i.e., the gerundive anchor “*triggering*” is transformed into “*triggered*” since it governs the open clausal complement of the anchor “*disputed*”. Algorithm 2 summarizes the event-wise simplification scheme.

Table 2: Simplification example

Original
“ <i>Baset al-Megrahi, the Libyan intelligence officer who was convicted in the 1988 Lockerbie bombing has died at his home in Tripoli, nearly three years after he was released from a Scottish prison.</i> ”
Sentence-wise simplification
“ <i>Baset al-Megrahi was convicted in the 1988 Lockerbie bombing has died at his home after he was released from a Scottish prison.</i> ”
Event-wise simplification
“ <i>Baset al-Megrahi was convicted in the 1988 Lockerbie bombing. Baset al-Megrahi has died at his home. He was released from a Scottish prison.</i> ”
Event-wise with pron. anaphora resolution
“ <i>Baset al-Megrahi was convicted in the 1988 Lockerbie bombing. Baset al-Megrahi has died at his home. Baset al-Megrahi was released from a Scottish prison.</i> ”

It has been shown that anaphoric mentions cause difficulties for people with cognitive disabilities (Ehrlich et al., 1999; Shapiro and Milkes, 2004). To investigate this phenomenon, we additionally employ *pronominal anaphora resolution* on top of event-wise simplification scheme. To resolve reference of anaphoric pronouns, we use the coreference resolution tool from Stanford Core NLP (Lee et al., 2011). An example of the original text snippet accompanied by its (1) sentence-wise simplification, (2) event-wise simplification, and (3) event-wise simplification with anaphoric pronoun resolution is given in Table 2.

4 Evaluation

The text is well-simplified if its readability is increased, while its grammaticality (syntactic correctness), meaning, and information relevance (semantic correctness) are preserved.

We measure the readability of the simplified text automatically with two commonly used formulae. However, we rely on human assessment of grammaticality and relevance, given that these aspects are difficult to capture automatically (Wubben et al., 2012). We employ a syntactically motivated baseline that retains only the main clause of a sentence and discards all subordinate clauses. We used Stanford constituency parser (Klein and Manning, 2003) to identify the main and subordinate clauses.

Readability. We collected 100 news stories from EMM NewsBrief,² an online news clustering ser-

²<http://emm.newsbrief.eu/NewsBrief/clusteredition/en/latest.html>

Table 3: Readability evaluation

Original vs.	KFL	SMOG	SL	DL	NS
Baseline	-27.7% ± 12.5%	-14.0% ± 8.0%	-38.5% ± 12.1%	-38.5% ± 12.1%	0.0% ± 0.0%
Sentence-wise	-30.1% ± 13.9%	-16.3% ± 9.2%	-44.3% ± 11.1%	-49.8% ± 11.5%	-9.9% ± 8.7%
Event-wise	-50.3% ± 12.6%	-30.8% ± 10.5%	-65.5% ± 9.3%	-63.4% ± 12.6%	-10.0% ± 39.7%
Pronom. anaphora	-47.8% ± 13.9%	-29.4% ± 10.6%	-63.6% ± 10.3%	-61.2% ± 14.4%	-10.0% ± 39.7%

vice, and simplified them with the proposed simplification schemes. For each original story and its simplified versions, we compute two standard readability scores – Kincaid-Flesch Grade Level (KFL) (Kincaid et al., 1975) and SMOG Index (McLaughlin, 1969). We also compute common-sense indicators of readability: average sentence length (SL), average document length (DL), and number of sentences (NS). Readability scores, relative to the readability of the original text and averaged over 100 news stories, are given in Table 3.

Event-wise simplification significantly ($p < 0.01$)³ increases the readability for all measures except NS. Large variance in NS for event-wise simplification is caused by large variance in number of factual events per news story. Descriptive news stories (e.g., political overviews) contain more sentences without any factual events, while sentences from factual stories (e.g., murders, protests) often contain several factual events, forming multiple sentences in the simplified text. Event-wise simplified texts are also significantly more readable than sentence-wise simplified texts ($p < 0.01$) for all measures except NS.

Human Evaluation. Readability scores provide no information about the content of the simplified text. In line with previous work on text simplification (Knight and Marcu, 2002; Woodsend and Lapata, 2011; Wubben et al., 2012; Drndarevic et al., 2013), we let human evaluators judge the grammaticality and content relevance of simplified text. Due to cognitive effort required for the annotation task we asked annotators to compare text snippets (consisting of a single sentence or two adjacent sentences) instead of whole news stories. For each simplification, evaluators were instructed to compare it with the respective original snippet and assign three different scores:

(1) *Grammaticality score* denotes the grammatical well-formedness of text on a 1-3 scale, where 1

denotes significant ungrammaticalities (e.g., missing subject or object as in “*Was prevented by the Chinese surveillance craft.*”), 2 indicates smaller grammatical inconsistencies (e.g., missing conjunctions or prepositions, as in “*Vessels blocked the arrest Chinese fishermen in disputed waters*”), and 3 indicates grammatical correctness;

(2) *Meaning score* denotes the degree to which relevant information from the original text is preserved semantically unchanged in the simplified text on a 1-3 scale, where 1 indicates that the most relevant information has not been preserved in its original meaning (e.g., “*Russians are tiring of Putin*” → “*Russians are tiring Putin*”), 2 denotes that relevant information is partially missing from the simplified text (e.g., “*Their daughter has been murdered and another daughter seriously injured.*” → “*Their daughter has been murdered.*”), and 3 means that all relevant information has been preserved;

(3) *Simplicity score* indicates the degree to which irrelevant information has been eliminated from the simplified text on a 1-3 scale, where 1 means that a lot of irrelevant information has been retained in the simplified text (e.g., “*The president, acting as commander in chief, landed in Afghanistan on Tuesday afternoon for an unannounced visit to the war zone*”), 2 denotes that some of the irrelevant information has been eliminated, but not all of it (e.g., “*The president landed in Afghanistan on Tuesday afternoon for an unannounced visit*”), and 3 indicates that only the most relevant information has been retained in the simplified text (e.g., “*The president landed in Afghanistan on Tuesday*”).

Note that *Meaning* and *Simplicity* can, respectively, be interpreted as recall and precision of information relevance. The less relevant information is preserved (i.e., false negatives), the lower the *Meaning* score will be. Similarly, the more irrelevant information is preserved (i.e., false positives), the lower the *Simplicity* score will be. Considering that the well-performing simplification method should both preserve relevant and eliminate irrelevant information, for each simplified text we com-

³2-tailed t-test if both samples are approx. normally distributed; Wilcoxon signed-rank test otherwise

Table 4: IAA for human evaluation

Aspect	weighted κ	Pearson	MAE
Grammaticality	0.68	0.77	0.18
Meaning	0.53	0.67	0.37
Simplicity	0.54	0.60	0.28

Table 5: Grammaticality and relevance

Scheme	Gramm. (1-3)	Relevance (1-3)
Baseline	2.57 ± 0.79	1.90 ± 0.64
Sentence-wise	1.98 ± 0.80	2.12 ± 0.61
Event-wise	2.70 ± 0.52	2.30 ± 0.54
Pronom. anaphora	2.68 ± 0.56	2.39 ± 0.57

pute *Relevance* score as the harmonic mean of its *Meaning* score and *Simplicity* score.

We compiled a dataset of 70 snippets of newspaper texts, each consisting of one or two sentences.⁴ We simplified these 70 snippets using the two proposed simplification schemes (and additional pronominal anaphora resolution) and the baseline, obtaining that way four different simplifications per snippet, i.e., 280 pairs of original and simplified text altogether. Three evaluators independently annotated the same 40 pairs on which we measured the inter-annotator agreement (IAA). Since we observed fair agreement, the evaluators proceeded by annotating the remaining 240 pairs. Pairwise averaged IAA in terms of three complementary metrics – Weighted Cohen’s Kappa (κ), Pearson correlation, and Mean Absolute Error (MAE) – is given in Table 4. As expected, IAA shows that grammaticality is less susceptible to individual interpretations than information (ir)relevance (i.e., *Meaning* and *Simplicity*). Nonetheless, we observe fair agreement for *Meaning* and *Simplicity* as well ($\kappa > 0.5$).

Finally, we evaluate the performance of the proposed simplification schemes on the 70 news snippets in terms of *Grammaticality* and *Relevance*. The results are shown in Table 5. All simplification schemes produce text which is significantly more relevant than the baseline simplification ($p < 0.05$ for sentence-wise scheme; $p < 0.01$ for the event-wise and pronominal anaphora schemes). However, sentence-wise simplification produces text which is significantly less grammatical than baseline simplification. This is because conjunctions and prepositions are often missing from sentence-wise simplifi-

cations as they do not form any event mention. The same issue does not arise in event-wise simplifications where each mention is converted into its own sentence, in which case eliminating conjunctions is grammatically desirable. Event-wise and pronominal anaphora schemes significantly outperform the sentence-wise simplification ($p < 0.01$) on both grammaticality and relevance. Most mistakes in event-wise simplifications originate from change of meaning caused by the incorrect extraction of event arguments (e.g., “*Nearly 3,000 soldiers have been killed in Afghanistan since the Taliban were ousted in 2001.*” \rightarrow “*Nearly 3,000 soldiers have been killed in Afghanistan in 2001.*”).

Overall, the event-wise scheme increases readability and produces grammatical text, preserving at the same time relevant content and reducing irrelevant content. Combined, experimental results for readability, grammaticality, and information relevance suggest that the proposed event-wise scheme is very suitable for text simplification.

5 Conclusion

Acknowledging that news stories are difficult to comprehend for people with reading disabilities, as well as the fact that events represent the most relevant information in news, we presented an event-centered approach to simplification of news. We identify factual event mentions with the state-of-the-art event extraction system and discard text that is not part of any of the factual events. Our experiments show that the event-wise simplification, in which factual events are converted to separate sentences, increases readability and retains grammaticality of the text, while preserving relevant information and discarding irrelevant information.

In future work we will combine event-based schemes with methods for lexical simplification. We will also investigate the effects of temporal ordering of events on text simplification, as texts with linear timelines are easier to follow. We also intend to employ similar event-based strategies for text summarization, given the notable similarities between text simplification and summarization.

Acknowledgments. This work has been partially supported by the Ministry of Science, Education and Sports, Republic of Croatia under the Grant 036-1300646-1986. We thank the PhD students from the University of Wolverhampton for their annotation efforts. We also thank the reviewers for their constructive comments.

⁴The dataset is freely available at <http://takelab.fer.hr/evsimplify>

References

- ACE. 2005. Evaluation of the detection and recognition of ACE: Entities, values, temporal expressions, relations, and events.
- D. Ahn. 2006. The stages of event extraction. In *Proceedings of the COLING/ACL 2006 Workshop on Annotating and Reasoning about Time and Events*, pages 1–8.
- S. M. Aluísio, L. Specia, T. A. S. Pardo, E. G. Maziero, and R. P. M. Fortes. 2008. Towards Brazilian Portuguese automatic text simplification systems. In *Proceedings of the eighth ACM symposium on Document engineering*, DocEng '08, pages 240–248, New York, NY, USA. ACM.
- G. Barlacchi and S. Tonelli. 2013. ERNESTA: A sentence simplification tool for childrens stories in italian. In *Computational Linguistics and Intelligent Text Processing*, pages 476–487. Springer.
- S. Bethard. 2008. *Finding Event, Temporal and Causal Structure in Text: A Machine Learning Approach*. Ph.D. thesis.
- J. Burstein, J. Shore, J. Sabatini, Y.W. Lee, and M. Ventura. 2007. The automated text adaptation tool. In *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, NAACL-Demonstrations '07, pages 3–4, Stroudsburg, PA, USA. Association for Computational Linguistics.
- B. Carretti, C. Belacchi, and C. Cornoldi. 2010. Difficulties in working memory updating in individuals with intellectual disability. *Journal of Intellectual Disability Research*, 54(4):337–345.
- J. Carroll, G. Minnen, Y. Canning, S. Devlin, and J. Tait. 1998. Practical simplification of English newspaper text to assist aphasic readers. In *Proceedings of AAAI-98 Workshop on Integrating Artificial Intelligence and Assistive Technology (1998)*, pp. 7–10 Key: citeulike:8717999, pages 7–10.
- J. Carroll, G. Minnen, D. Pearce, Y. Canning, S. Devlin, and J. Tait. 1999. Simplifying text for language-impaired readers. In *Proceedings of the 9th Conference of the European Chapter of the ACL (EACL'99)*, pages 269–270.
- A. X. Chang and C. D. Manning. 2012. SUTIME: A library for recognizing and normalizing time expressions. *Language Resources and Evaluation*.
- S. Devlin and G. Unthank. 2006. Helping aphasic people process online information. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility*, Assets '06, pages 225–226, New York, NY, USA. ACM.
- S. Devlin. 1999. *Simplifying Natural Language Text for Aphasic Readers*. Ph.D. thesis, University of Sunderland, UK.
- B. Drndarevic, S. Štajner, S. Bott, S. Bautista, and H. Saggion. 2013. Automatic text simplification in Spanish: A comparative evaluation of complementing components. In *Proceedings of the 12th International Conference on Intelligent Text Processing and Computational Linguistics. Lecture Notes in Computer Science. Samos, Greece, 24-30 March, 2013.*, pages 488–500.
- M. Ehrlich, M. Remond, and H. Tardieu. 1999. Processing of anaphoric devices in young skilled and less skilled comprehenders: Differences in metacognitive monitoring. *Reading and Writing*, 11(1):29–63.
- L. Feng. 2009. Automatic readability assessment for people with intellectual disabilities. In *SIGACCESS Access. Comput.*, number 93, pages 84–91. ACM, New York, NY, USA, jan.
- J. Finkel, T. Grenager, and C. Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 363–370. Association for Computational Linguistics.
- G. Freyhoff, G. Hess, L. Kerr, B. Tronbacke, and K. Van Der Veken, 1998. *Make it Simple, European Guidelines for the Production of Easy-to-Read Information for People with Learning Disability*. ILSMH European Association, Brussels.
- Goran Glavaš and Jan Šnajder. 2013. Exploring coreference uncertainty of generically extracted event mentions. In *Proceedings of 14th International Conference on Intelligent Text Processing and Computational Linguistics*, pages 408–422. Springer.
- C. Grover, R. Tobin, B. Alex, and K. Byrne. 2010. Edinburgh-LTG: TempEval-2 system description. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 333–336. Association for Computational Linguistics.
- K. Inui, A. Fujita, T. Takahashi, R. Iida, and T. Iwakura. 2003. Text simplification for reading assistance: A project note. In *Proceedings of the second international workshop on Paraphrasing - Volume 16, PARAPHRASE '03*, pages 9–16, Stroudsburg, PA, USA. Association for Computational Linguistics.
- J. P. Kincaid, R. P. Fishburne Jr, R. L. Rogers, and B. S. Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, DTIC Document.
- D. Klein and C. D. Manning. 2003. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 423–430. Association for Computational Linguistics.

- K. Knight and D. Marcu. 2002. Summarization beyond sentence extraction: A probabilistic approach to sentence compression. *Artificial Intelligence*, 139:91–107.
- P. Lal and S. Ruger. 2002. Extract-based Summarization with Simplification. In *Proceedings of the ACL 2002 Automatic Summarization / DUC 2002 Workshop*.
- H. Lee, Y. Peirsman, A. Chang, N. Chambers, M. Surdeanu, and D. Jurafsky. 2011. Stanford’s multi-pass sieve coreference resolution system at the CoNLL-2011 shared task. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task*, pages 28–34. Association for Computational Linguistics.
- H. Llorens, E. Saquete, and B. Navarro. 2010. Tipsem (english and spanish): Evaluating CRFs and semantic roles in TempEval-2. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 284–291. Association for Computational Linguistics.
- G. H. McLaughlin. 1969. SMOG grading: A new readability formula. *Journal of reading*, 12(8):639–646.
- C. Orasan, Evans. R., and I. Dornescu, 2013. *Towards Multilingual Europe 2020: A Romanian Perspective*, chapter Text Simplification for People with Autistic Spectrum Disorders, pages 287–312. Romanian Academy Publishing House, Bucharest.
- Z. Pan and G. M. Kosicki. 1993. Framing analysis: An approach to news discourse. *Political communication*, 10(1):55–75.
- H. Pimperton and K. Nation. 2010. Suppressing irrelevant information from working memory: Evidence for domain-specific deficits in poor comprehenders. *Journal of Memory and Language*, 62(4):380–391.
- J. Pustejovsky, J. Castano, R. Ingria, R. Sauri, R. Gaizauskas, A. Setzer, G. Katz, and D. Radev. 2003a. TimeML: Robust specification of event and temporal expressions in text. *New Directions in Question Answering*, 2003:28–34.
- J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizauskas, A. Setzer, D. Radev, B. Sundheim, D. Day, L. Ferro, et al. 2003b. The TimeBank corpus. In *Corpus Linguistics*, volume 2003, page 40.
- H. Saggion, E. Gómez Martínez, E. Etayo, A. Anula, and L. Bourg. 2011. Text simplification in Simplext: Making text more accessible. *Revista de la Sociedad Española para el Procesamiento del Lenguaje Natural*.
- C. Scarton, M. de Oliveira, A. Candido, C. Gasperin, and S. M. Alusio. 2010. SIMPLIFICA: A tool for authoring simplified texts in Brazilian Portuguese guided by readability assessments. In *Proceedings of the NAACL HLT 2010: Demonstration Session*, pages 41–44.
- A. Shapiro and A. Milkes. 2004. Skilled readers make better use of anaphora: A study of the repeated-name penalty on text comprehension. *Electronic Journal of Research in Educational Psychology*, 2(2):161–180.
- T. A. Van Dijk. 1985. Structures of news in the press. *Discourse and communication: New approaches to the analysis of mass media discourse and communication*, 10:69.
- M. Verhagen, R. Sauri, T. Caselli, and J. Pustejovsky. 2010. Semeval-2010 task 13: TempEval-2. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 57–62. Association for Computational Linguistics.
- S. Štajner, R. Evans, C. Orasan, and R. Mitkov. 2012. What Can Readability Measures Really Tell Us About Text Complexity? In *Proceedings of the LREC’12 Workshop: Natural Language Processing for Improving Textual Accessibility (NLP4ITA)*, Istanbul, Turkey.
- K. Woodsend and M. Lapata. 2011. Learning to simplify sentences with quasi-synchronous grammar and integer programming. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Z. Wu and M. Palmer. 1994. Verbs semantics and lexical selection. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pages 133–138. Association for Computational Linguistics.
- S. Wubben, A. van den Bosch, and E. Krahrmer. 2012. Sentence simplification by monolingual machine translation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers - Volume 1, ACL ’12*, pages 1015–1024, Stroudsburg, PA, USA. Association for Computational Linguistics.