

A Readable Read: Automatic Assessment of Language Learning Materials based on Linguistic Complexity

ILDIKÓ PILÁN¹, SOWMYA VAJJALA², AND ELENA VOLODINA¹

¹ *University of Gothenburg, Sweden*

² *Universität Tübingen, Germany*

ABSTRACT

Corpora and web texts can become a rich language learning resource if we have a means of assessing whether they are linguistically appropriate for learners at a given proficiency level. In this paper, we aim at addressing this issue by presenting the first approach for predicting linguistic complexity for Swedish second language learning material on a 5-point scale. After showing that the traditional Swedish readability measure, Läsbarhetsindex (LIX), is not suitable for this task, we propose a supervised machine learning model, based on a range of linguistic features, that can reliably classify texts according to their difficulty level. Our model obtained an accuracy of 81.3% and an F-score of 0.8, which is comparable to the state of the art in English and is considerably higher than previously reported results for other languages. We further studied the utility of our features with single sentences instead of full texts since sentences are a common linguistic unit in language learning exercises. We trained a separate model on sentence-level data with five classes, which yielded 63.4% accuracy. Although this is lower than the document level performance, we achieved an adjacent accuracy of 92%. Furthermore, we found that using a combination of different features, compared to using lexical features alone, resulted in 7% improvement in classification accuracy at the sentence level, whereas at the document level, lexical features were more dominant. Our models

**This is a pre-print version of the paper, before proper
formatting and copyediting by the editorial staff.**

are intended for use in a freely accessible web-based language learning platform for the automatic generation of exercises.

KEYWORDS: *readability, linguistic complexity, machine learning, computer-assisted language learning, CEFR*

1 INTRODUCTION

Linguistic information provided by Natural Language Processing (NLP) tools has good potential for turning the continuously growing amount of digital text into interactive and personalized language learning material. Our work aims at overcoming one of the fundamental obstacles in this domain of research, namely how to assess the linguistic complexity of texts and sentences from the perspective of second and foreign language (L2) learners.

There are a number of readability models relying on NLP tools to predict the difficulty (readability) level of a text [1–6]. The linguistic features explored so far for this task incorporate information, among others, from part-of-speech (POS) taggers and dependency parsers. Cognitively motivated features have also been proposed, for example, in the Coh-Metrix [3]. Although the majority of previous work focuses primarily on document-level analysis, a finer-grained, sentence-level readability has received increasing interest in recent years [7–9].

The previously mentioned studies target mainly native language (L1) readers including people with low literacy levels or mild cognitive disabilities. Our focus, however, is on building a model for predicting the proficiency level of texts and sentences used in L2 teaching materials. This aspect has been explored for English [10–13], French [14], Portuguese [15] and, without the use of NLP, for Dutch [16].

Readability for the Swedish language has a rather long tradition. One of the most popular, easy-to-compute formulas is LIX (Läsbarhetsindex, ‘Readability index’) proposed in [17]. This measure combines the average number of words per sentence in the text with the percentage of long words, i.e. tokens consisting of more

than six characters. Besides traditional formulas, supervised machine learning approaches have also been tested. Swedish document-level readability with a native speaker focus is described in [5] and [18]. For L2 Swedish, only a binary sentence-level model exists [9], but comprehensive and highly accurate document- and sentence-level models for multiple proficiency levels have not been developed before.

In this paper, we present a machine learning model trained on course books currently in use in L2 Swedish classrooms. Our goal was to predict linguistic complexity of material written by teachers and course book writers for learners, rather than assessing learner-produced texts. We adopted the scale from the Common European Framework of Reference for Languages (CEFR) [19] which contains guidelines for the creation of teaching material and the assessment of L2 proficiency. CEFR proposes six levels of language proficiency: A1 (beginner), A2 (elementary), B1 (intermediate), B2 (upper intermediate), C1 (advanced) and C2 (proficient). Since sentences are a common unit in language exercises, but remain less explored in the readability literature, we also investigate the applicability of our approach to sentences, performing a 5-way classification (levels A1-C1). Our document-level model achieves a state-of-the-art performance (F-score of 0.8), however, there is room for improvement in sentence-level predictions. We plan to make our results available through the online intelligent computer-assisted language learning platform Lärka³, both as corpus-based exercises for teachers and learners of L2 Swedish and as web-services for researchers and developers.

In the following sections, we first describe our datasets (section 2) and features (section 3), then we present the details and the results of our experiments in section 4. Finally, section 5 concludes our work and outlines further directions of research within this area.

³ <http://spraakbanken.gu.se/larka/>

2 DATASETS

Our dataset is a subset of COCTAILL, a corpus of course books covering five CEFR levels (A1-C1) [20]. This corpus consists of twelve books (from four different publishers) whose usability and level have been confirmed by Swedish L2 teachers. The course books have been annotated both content-wise (e.g. exercises, lists) and linguistically (e.g. with POS and dependency tags) [20]. We collected a total of 867 texts (reading passages) from this corpus. We excluded texts that are primarily based on dialogues from the current experiments due to their specific linguistic structure, with the aim of scaling down differences connected to text genres rather than linguistic complexity. We plan to study the readability of dialogues and compare them to non-dialogue texts in the future.

Besides reading passages, i.e. texts, the COCTAILL corpus contains a number of sentences independent from each other, i.e. not forming a coherent text, in the form of lists of sentences and *language examples*. This latter category consists of sentences illustrating the use of specific grammatical patterns or lexical items. Collecting these sentences, we built a sentence-level dataset consisting of 1874 instances. The information encoded in the content-level annotation of COCTAILL (XML tags *list*, *language_example* and the attribute *unit*) enabled us to include only complete sentences and exclude sentences containing gaps and units larger or smaller than a sentence (e.g. texts, phrases, single words etc.). The CEFR level of both sentences and texts has been derived from the CEFR level of the lesson (chapter) they appeared in. In Table 1, columns 2-5 give an overview of the distribution of texts across levels and their mean length in sentences.⁴ The distribution of sentences per level is presented in the last two columns of Table 1. COCTAILL contained a somewhat more limited amount of B2 and C1 level sentences in the form of lists and language examples, pos-

⁴ The number of different books and publishers is reported per each level, some books spanning more levels.

sibly because learners handle larger linguistic units with more ease at higher proficiency levels.

Table 1. The distribution of items per CEFR level in the datasets.

CEFR	Document level				Sentence level	
	Books	Publ.	Texts	Mean nr. sent	Books	Sentences
A1	4	3	49	14.0	4	505
A2	4	3	157	13.8	4	754
B1	5	3	258	17.9	4	408
B2	4	3	288	26.6	3	124
C1	2	2	115	42.1	1	83
Total	12	4	867	-	4	1874

3 FEATURES

We developed our features based on information both from previous literature [10, 4, 14, 5, 9] and a grammar book for Swedish L2 learners [21]. The set of features can be divided in the following five subgroups: length-based, lexical, morphological, syntactic and semantic features (Table 2).

Length-based (LEN): These features include sentence length in number of tokens (#1) and characters (#4), extra-long words (longer than thirteen characters) and the traditional Swedish readability formula, LIX (see section 1). For the sentence-level analysis, instead of the ratio of number of tokens to the number of sentences in the text, we considered the number of tokens in one sentence.

Lexical (LEX): Similar to [9], we used information from the Kelly list [22], a lexical resource providing a CEFR level and frequencies per lemma based on a corpus of web texts. Thus, this word list is entirely independent from our dataset. Instead of percentages, we used *incidence scores* (INCSC) per 1000 words to reduce the influence of sentence length on feature values. The INCSC of a category was computed as 1000 divided by the number of tokens in the text or sentence multiplied by the count of the category

in the sentence. We calculated the INCS of words belonging to each CEFR level (#6 - #11). In features #12 and #13 we considered *difficult* all tokens whose level was above the CEFR level of the text or sentence. We computed also the INCS of tokens not present in the Kelly list (#14), tokens for which the lemmatizer did not find a corresponding lemma form (#15), as well as average log frequencies (#16).

Morphological (MORPH): We included the variation (the ratio of a category to the ratio of lexical tokens - i.e. nouns, verbs, adjectives and adverbs) and the INCS of all lexical categories together with the INCS of punctuations, particles, sub- and conjunctions (#34, #51). Some additional features, using insights from L2 teaching material [21], captured fine-grained inflectional information such as the INCS of neuter gender nouns and the ratio of different verb forms to all verbs (#52 - #56). Instead of simple type-token ratio (TTR) we used a bilogarithmic and a square root TTR as in [4]. Moreover, nominal ratio [5], the ratio of pronouns to prepositions [14], and two *lexical density* features were also included: the ratio of lexical words to all non-lexical categories (#48) and to all tokens (#49). Relative structures (#57) consisted of relative adverbs, determiners, pronouns and possessives.

Syntactic (SYNT): Some of these features were based on the length (depth) and the direction of dependency arcs⁵ (#17 - #21). We complemented this, among others, with the INCS of relative clauses in clefts⁶ (#26), and the INCS of pre- and postmodifiers (e.g. adjectives and prepositional phrases) [5].

Semantic (SEM): Features based on information from SALDO [24], a Swedish lexical-semantic resource. We used the average number of senses per token as in [9] and included also the average number of noun senses per nouns. Once reliable word-sense disambiguation methods become available for Swedish, additional

⁵ The tags were obtained with the MaltParser [23].

⁶ Sentences that begin with a constituent receiving particular focus, followed by a relative clause. E.g.: It is John (whom) Jack is waiting for.

features based on word senses could be taken into consideration here.

The complete set of 61 features is presented in Table 2. Throughout this paper we will refer to the machine learning models using this set of features, unless otherwise specified. Features for both document- and sentence-level analyses were extracted for each sentence, the values being averaged over all sentences in the text in the document-level experiments to ensure comparability.

4 EXPERIMENTS AND RESULTS

4.1 *Experimental Setup*

We explored different classification algorithms for this task using the machine learning toolkit WEKA [25]. These included: (1) a multinomial logistic regression model with ridge estimator, (2) a multilayer perceptron, (3) a support vector machine learner, Sequential Minimal Optimization (SMO), and (4) a decision tree (J48). For each of these, the default parameter settings have been used as implemented in WEKA.

We considered classification accuracy, F-score and Root Mean Squared Error (RMSE) as evaluation measures for our approach. We also included a confusion matrix, as we deal with a dataset that is unbalanced across CEFR levels. The scores were obtained by performing a ten-fold Cross-Validation (CV).

4.2 *Document-Level Experiments*

We trained document-level classification models, comparing the performance between different subgroups of features. We had two baselines: a majority classifier (MAJORITY), with B2 as majority class, and the LIX readability score. Table 3 shows the type of subgroup (*Type*), the number of features (*Nr*) and three evaluation metrics using logistic regression.

Not only was accuracy very low with LIX, but this measure also classified 91.6% of the instances as B2 level. Length-based,

Table 2. The complete feature set.

Nr.	Feature Name	Nr.	Feature Name
<i>Length-based</i>		<i>Morphological</i>	
1	Sentence length	30	Modal verbs to verbs
2	Average token length	31	Particle INCSC
3	Extra-long words	32	3SG pronoun INCSC
4	Number of characters	33	Punctuation INCSC
5	LIX	34	Subjunction INCSC
<i>Lexical</i>		35	S-verb INCSC
6	A1 lemma INCSC	36	S-verbs to verbs
7	A2 lemma INCSC	37	Adjective INCSC
8	B1 lemma INCSC	38	Adjective variation
9	B2 lemma INCSC	39	Adverb INCSC
10	C1 lemma INCSC	40	Adverb variation
11	C2 lemma INCSC	41	Noun INCSC
12	Difficult word INCSC	42	Noun variation
13	Difficult noun and verb INCSC	43	Verb INCSC
14	Out-of-Kelly INCSC	44	Verb variation
15	Missing lemma form INCSC	45	Nominal ratio
16	Avg. Kelly log frequency	46	Nouns to verbs
<i>Syntactic</i>		47	Function word INCSC
17	Average dependency length	48	Lexical words to non-lexical words
18	Dependency arcs longer than 5	49	Lexical words to all tokens
19	Longest dependency from root node	50	Neuter gender noun INCSC
20	Ratio of right dependency arcs	51	Con- and subjunction INCSC
21	Ratio of left dependency arcs	52	Past participles to verbs
22	Modifier variation	53	Present participles to verbs
23	Pre-modifier INCSC	54	Past verbs to verbs
24	Post-modifier INCSC	55	Present verbs to verbs
25	Subordinate INCSC	56	Supine verbs to verbs
26	Relative clause INCSC	57	Relative structure INCSC
27	Prepositional complement INCSC	58	Bilog type-token ratio
<i>Semantic</i>		59	Square root type-token ratio
28	Avg. nr. of senses per token	60	Pronouns to nouns
29	Noun senses per noun	61	Pronouns to prepositions

semantic and syntactic features in isolation showed similar or only slightly better performance than the baselines, therefore we excluded them from Table 3. Lexical features, however, had a strong discriminatory power without an increase in bias towards the majority classes. Using this subset of features only, we achieved ap-

Table 3. Document-level classification results.

Type	Nr	Acc (%)	F	RMSE
MAJORITY	-	33.2	0.17	0.52
LIX	1	34.9	0.22	0.38
LEX	11	80.3	0.80	0.24
ALL	61	81.3	0.81	0.27

proximately the same performance (0.8 F) as with the complete set of features, ALL (0.81 F). This suggests that lexical information alone can successfully distinguish the CEFR level of course book texts at the document level. Using the complete feature set we obtained 81% accuracy and 97% *adjacent accuracy* (when misclassifications to adjacent classes are considered correct). The same scores with lexical features (LEX) only were 80.3% (accuracy) and 98% (adjacent accuracy).

Accuracy scores using other learning algorithms were significantly lower (see Table 4), therefore, we report only the results of the logistic regression classifier in the subsequent sections.

Table 4. Accuracy scores (in %) for other learning algorithms.

Type	Nr	Perceptron	SMO	J48
LEX	11	77.4	42.1	55
ALL	61	62.2	52.7	50.5

Instead of classification, some readability studies (e.g. [11, 15]) employed linear regression for this task. For a better comparability, we applied also a linear regression model to our data which yielded a correlation of 0.8 and an RMSE of 0.65.

To make sure that our system was not biased towards the majority classes B1 and B2, we inspected the confusion matrix (Table 5) after classification using ALL. We can observe from Table 5 that the system performs better at A1 and C1 levels, where confusion occurred only with adjacent classes. Similar to the findings in [14] for French, classes in the middle of the scale were harder to distin-

guish. Most misclassifications in our material occurred at A2 level (23%) followed by B1 and B2 level, (20% and 17% respectively).

Table 5. Confusion matrix for feature set ALL at document level.

		Predictions						
		A1	A2	B1	B2	C1		
37	12	0	0	0	0	A1	L	
12	121	18	5	1	1	A2	a	
4	11	206	24	13	13	B1	b	
0	5	21	238	24	24	B2	e	
0	0	0	12	103	103	C1	l	

To establish the external validity of our approach, we tested it on a subset of LÄSBART [5], a corpus of Swedish easy-to-read (ETR) texts previously employed for Swedish L1 readability studies [5, 18]. We used 18 fiction texts written for children between ages nine to twelve, half of which belonged to the ETR category and the rest were unsimplified. Our model generalized well to unseen data, it classified all ETR texts as B1 and all ordinary texts as C1 level, thus correctly identifying in all cases the relative difference in complexity between the documents of the two categories.

Although a direct comparison with other studies is difficult because of the target language, the nature of the datasets and the number of classes used, in terms of absolute numbers, our model achieves comparable performance with the state-of-the-art systems for English[10, 13]. Other studies for non-English languages using CEFR levels include: [14], reporting 49.1% accuracy for a French system distinguishing six classes; and [15] achieving 29.7% accuracy on a smaller Portuguese dataset with five levels.

4.3 Sentence-Level Experiments

After building good classification models at document level, we explored the usability of our approach at the sentence level. Sentences are particularly useful in Computer-Assisted Language Learning

(CALL) applications, among others, for generating sentence-based multiple choice exercises, e.g. [26], or vocabulary examples [27]. Furthermore, multi-class readability classification of sentence-level material intended for second language learners has not been previously investigated in the literature.

With the same methodology (section 4.1) and feature set (section 3) used at the document level, we trained and tested classification models based on the sentence-level data (see section 2). The results are shown in Table 6.

Table 6. Sentence-level classification results.

Type	Nr	Acc (%)	F	RMSE
MAJORITY	-	40.2	0.23	0.49
LIX	1	41.4	0.3	0.38
LEX	11	56.8	0.53	0.33
ALL	61	63.4	0.63	0.31

Although the majority baseline in the case of sentences was 7% higher than the one for texts (Table 3), the classification accuracy for sentences using all features was only 63.4%. This is a considerable drop (−18%) in performance compared to the document level (81.3% accuracy). It is possible that the features did not capture differences between the sentences because the amount of context is more limited on the fine-grained level. It is interesting to note that, although there was no substantial performance difference between LEX and ALL at a document level, the model with all the features performed 7% better at sentence level.

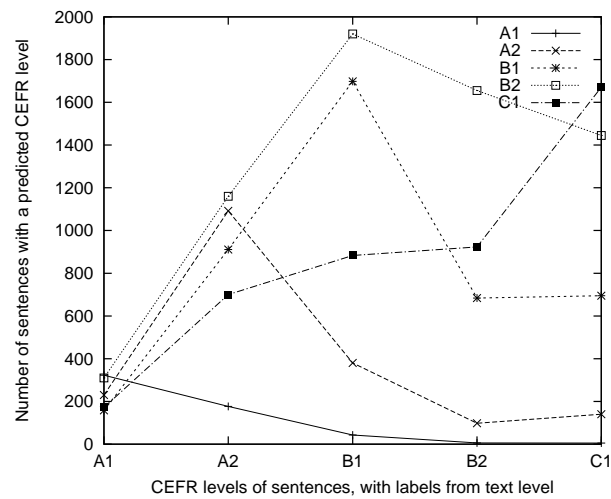
Most misclassifications occurred, however, within a distance of one class only, thus the adjacent accuracy of the sentence-level model was still high, 92% (see Table 7). Predictions were noticeably more accurate for classes A1, A2 and B1 which had a larger number of instances.

In the next step, we applied the sentence-level model on the document-level data to explore how homogeneous texts were in terms of the CEFR level of the sentences they contained. Figure 1

Table 7. Confusion matrix for feature set ALL at sentence level.

					Predictions	
A1	A2	B1	B2	C1	A1	L
371	123	9	2	0	A1	L
120	541	78	11	4	A2	a
27	136	212	23	10	B1	b
8	34	39	30	13	B2	e
0	18	21	9	35	C1	l

shows that texts at each CEFR level contain a substantial amount of sentences of the same level of the whole text, but they also include sentences classified as belonging to other CEFR levels.

**Fig. 1.** Distribution of sentences per CEFR level in the document-level data.

Finally, as in the case of the document-level analysis, we tested our sentence-level model also on an independent dataset (SENREAD), a small corpus of sentences with gold-standard CEFR annotation. This data was created during a user-based evaluation study [28]

and it consists of 196 sentences from generic corpora, i.e. originally not L2 learner-focused corpora, rated as being suitable at B1 or being at a level higher than B1. We used this corpus along with the judgments of the three participating teachers. Since SENREAD had only two categories - $\leq B1$ and $> B1$, we combined the model's predictions into two classes - A1, A2, B1 were considered as $\leq B1$ and B2, C1 were considered as $> B1$. The majority baseline for the dataset was 65%, $\leq B1$ being the class with most instances. The model trained on COCTAILL sentences predicted with 73% accuracy teachers' judgments, an 8% improvement over the majority baseline. There was a considerable difference between the precision score of the two classes, which was 85.4% for $\leq B1$, and only 48.5% for $> B1$.

Previously published results on sentence-level data include [7], who report 66% accuracy for a binary classification task for English and [8] who obtained an accuracy between 78.9% and 83.7% for Italian binary class data using different kinds of datasets. Neither of these studies, however, had a non-native speaker focus. [9] report 71% accuracy for Swedish binary sentence-level classification from an L2 point of view. Both the adjacent accuracy of our sentence-level model (92%) and the accuracy score obtained with that model on SENREAD (73%) improve on that score. It is also worth mentioning that the labels in the dataset from [9] were based on the assumption that all sentences in a text belong to the same difficulty level which, being an approximation (as also Figure 1 shows), introduced some noise in that data.

Although more analysis would be needed to refine the sentence-level model, our current results indicate that a rich feature set that considers multiple linguistic dimensions may result in an improved performance. In the future, the dataset could be expanded with more gold-standard sentences, which may improve accuracy. Furthermore, an interesting direction to pursue would be to verify whether providing finer-grained readability judgments is a more challenging task also for human raters.

5 CONCLUSION AND FUTURE WORK

We proposed an approach to assess the proficiency (CEFR) level of Swedish L2 course book texts based on a variety of features. Our document-level model, the first for L2 Swedish, achieved an F-score of 0.8, hence, it can reliably distinguish between proficiency levels. Compared to the wide-spread readability measure for Swedish, LIX, we achieved a substantial gain in terms of both accuracy and F-score (46% and 0.6 higher respectively). The accuracy of the sentence-level model remained lower than that of the text-level model, nevertheless, using the complete feature set the system performed 23% and 22% above the majority baseline and LIX respectively. Misclassifications of more than one level did not occur in more than 8% of sentences, thus, in terms of adjacent accuracy, our sentence-level model improved on previous results for L2 Swedish readability [9].

Most notably, we have found that taking into consideration multiple linguistic dimensions when assessing linguistic complexity is especially useful for sentence-level analysis. In our experiments, using only word-frequency features was almost as predictive as a combination of all features for the document level, but the latter made more accurate predictions for sentences, resulting in a 7% difference in accuracy. Besides L2 course book materials, we tested both our document- and sentence-level models also on unseen data with promising results.

In the future, a more detailed investigation is needed to understand the performance drop between document and sentence level. Acquiring more sentence-level annotated data and exploring new features relying on lexical-semantic resources for Swedish would be interesting directions to pursue. Furthermore, we intend to test the utility of this approach in a real-world web application involving language learners and teachers.

REFERENCES

1. Collins-Thompson, K., Callan, J.P.: A language modeling approach to predicting reading difficulty. In: HLT-NAACL. (2004) 193–200
2. Schwarm, S.E., Ostendorf, M.: Reading level assessment using support vector machines and statistical language models. In: Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, Association for Computational Linguistics (2005) 523–530
3. Graesser, A.C., McNamara, D.S., Kulikowich, J.M.: Coh-Metrix providing multilevel analyses of text characteristics. *Educational Researcher* **40**(5) (2011) 223–234
4. Vajjala, S., Meurers, D.: On improving the accuracy of readability classification using insights from second language acquisition. In: Proceedings of the Seventh Workshop on Innovative Use of NLP for Building Educational Applications, Association for Computational Linguistics (2012) 163–173
5. Heimann Mühlenbock, K.: I see what you mean. PhD thesis, University of Gothenburg (2013)
6. Collins-Thompson, K.: Computational assessment of text readability: A survey of current and future research. *Recent Advances in Automatic Readability Assessment and Text Simplification*. Special issue of *International Journal of Applied Linguistics* **6** (2014) 97–135
7. Vajjala, S., Meurers, D.: Assessing the relative reading level of sentence pairs for text simplification. In: Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL-14), Gothenburg, Sweden, Association for Computational Linguistics (2014)
8. Dell’Orletta, F., Wieling, M., Cimino, A., Venturi, G., Montemagni, S.: Assessing the readability of sentences: Which corpora and features? *ACL 2014* (2014) 163
9. Pilán, I., Volodina, E., Johansson, R.: Rule-based and machine learning approaches for second language sentence-level readability. In: Proceedings of the Ninth Workshop on Innovative Use of NLP for Building Educational Applications, Baltimore, Maryland, Association for Computational Linguistics (June 2014) 174–184
10. Heilman, M.J., Collins-Thompson, K., Callan, J., Eskenazi, M.: Combining lexical and grammatical features to improve readability measures for first and second language texts. In: Proceedings of NAACL HLT. (2007) 460–467
11. Huang, Y.T., Chang, H.P., Sun, Y., Chen, M.C.: A robust estimation scheme of reading difficulty for second language learners. In: *Advanced Learning Technologies (ICALT), 2011 11th IEEE International Conference on*, IEEE (2011) 58–62
12. Zhang, L., Liu, Z., Ni, J.: Feature-based assessment of text readability. In: *Internet Computing for Engineering and Science (ICICSE), 2013 Seventh International Conference on*, IEEE (2013) 51–54

13. Salesky, E., Shen, W.: Exploiting morphological, grammatical, and semantic correlates for improved text difficulty assessment. In: Proceedings of the Ninth Workshop on Innovative Use of NLP for Building Educational Applications, Baltimore, Maryland, Association for Computational Linguistics (June 2014) 155–162
14. François, T., Fairon, C.: An AI readability formula for French as a foreign language. In: Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, Association for Computational Linguistics (2012) 466–477
15. Branco, A., Rodrigues, J., Costa, F., Silva, J., Vaz, R.: Rolling out text categorization for language learning assessment supported by language technology. In: Computational Processing of the Portuguese Language. Springer (2014) 256–261
16. Velleman, E., van der Geest, T.: Online test tool to determine the CEFR reading comprehension level of text. *Procedia Computer Science* **27** (2014) 350–358
17. Björnsson, C.H.: *Läsbarhet. Liber* (1968)
18. Falkenjack, J., Heimann Mühlenbock, K., Jönsson, A.: Features indicating readability in Swedish text. In: Proceedings of the 19th Nordic Conference of Computational Linguistics (NODALIDA 2013). (2013) 27–40
19. Council of Europe: *Common European Framework of Reference for Languages: Learning, Teaching, Assessment*. Cambridge University Press (2001)
20. Volodina, E., Pilán, I., Eide, S.R., Heidarsson, H.: You get what you annotate: a pedagogically annotated corpus of coursebooks for Swedish as a Second Language. *NEALT Proceedings Series Vol. 22* (2014) 128
21. Fasth, C., Kannermark, A.: *Form i fokus: övningsbok i svensk grammatik. Del B*. Folkuniv. Förlag, Lund (1997)
22. Volodina, E., Kokkinakis, S.J.: Introducing the Swedish Kelly-list, a new lexical e-resource for Swedish. In: Proceedings of LREC. (2012) 1040–1046
23. Nivre, J., Hall, J., Nilsson, J., Chanev, A., Eryigit, G., Kübler, S., Marinov, S., Marsi, E.: MaltParser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering* **13**(02) (2007) 95–135
24. Borin, L., Forsberg, M., Lönngrén, L.: SALDO: a touch of yin to WordNet’s yang. *Language Resources and Evaluation* **47**(4) (2013) 1191–1211
25. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA data mining software: An update. In: *The SIGKDD Explorations*. Volume 11. (2009) 10–18
26. Volodina, E., Pilán, I., Borin, L., Tiedemann, T.L.: A flexible language learning platform based on language resources and web services. In: Proceedings of LREC 2014, Reykjavik, Iceland (may 2014)
27. Segler, T.M.: Investigating the selection of example sentences for unknown target words in ICALL reading texts for L2 German. PhD thesis, University of Edinburgh (2007)

28. Pilán, I., Volodina, E., Johansson, R.: Automatic selection of suitable sentences for language learning exercises. In: 20 Years of EUROCALL: Learning from the Past, Looking to the Future. 2013 EUROCALL Conference, 11th to 14th September 2013 Évora, Portugal, Proceedings. (2013) 218–225

ILDIKÓ PILÁN

SWEDISH LANGUAGE BANK,
UNIVERSITY OF GOTHENBURG,
40530 GOTHENBURG, SWEDEN

E-MAIL: <ILDIKO.PILAN@SVENSKA.GU.SE>

SOWMYA VAJJALA

LEAD GRADUATE SCHOOL,
SEMINAR FÜR SPRACHWISSENSCHAFT
UNIVERSITÄT TÜBINGEN,
72074 TÜBINGEN, GERMANY

E-MAIL: <SOWMYA@SFS.UNI-TUEBINGEN.DE>

ELENA VOLODINA

SWEDISH LANGUAGE BANK,
UNIVERSITY OF GOTHENBURG,
40530 GOTHENBURG, SWEDEN

E-MAIL: <ELENA.VOLODINA@SVENSKA.GU.SE>