Lärka: From Language Learning Platform to Infrastructure for Research on Language Learning

David Alfter Lars Borin Ildikó Pilán
Språkbanken, University of Gothenburg, Sweden
david.alfter|lars.borin|ildiko.pilan@gu.se

Therese Lindström Tiedemann
Department of Finnish, Finno-Ugrian
and Scandinavian Studies
University of Helsinki, Finland
therese.lindstromtiedemann@helsinki.fi

Elena Volodina
Språkbanken
University of Gothenburg, Sweden
elena.volodina@gu.se

Abstract

Lärka is an Intelligent Computer-Assisted Language Learning (ICALL) platform developed at Språkbanken, as a flexible and a valuable source of additional learning material (e.g. via corpus-based exercises) and a support tool for both teachers and L2 learners of Swedish and students of (Swedish) linguistics. Nowadays, Lärka is being adapted into a building block in an emerging second language research infrastructure within a larger context of the text-based research infrastructure developed by the national Swedish Language bank, Språkbanken, and SWE-CLARIN.

Lärka has recently received a new responsive user interface adapted to different devices with different screen sizes. Moreover, the system has also been augmented with new functionalities. These recent additions aim at improving the usability and the usefulness of the platform for pedagogical purposes. The most important development, though, is the adaptation of the platform to serve as a component in an e-infrastructure supporting research on language learning and multilingualism. Thanks to Lärka's *service-oriented architecture*, most functionalities are also available as web services which can be easily re-used by other applications.

1 Introduction

Lärka[1] is an Intelligent Computer-Assisted Language Learning (ICALL) platform developed at the CLARIN B Center Språkbanken Text (University of Gothenburg, Sweden). Lärka development started in the project *A system architecture for ICALL* (Volodina et al., 2012), the initial goal being to re-implement a previous tool, ITG, used up until then for teaching grammar (Borin and Saxena, 2004) with modern technology. The new application, Lärka, gradually developed into a platform for language learning covering two groups of learners – second/foreign language learners of Swedish and students of (Swedish) linguistics. Lärka is an openly available web-based tool that builds on a variety of existing SWE-CLARIN language resources such as Korp (Borin et al., 2012) for querying corpora, Karp (Borin et al., 2013b) for querying lexical resources and language technology tools (Borin et al., 2017). Thanks to its service-oriented architecture, Lärka functionalities can be re-used in other applications (Volodina et al., 2014b).

In parallel to exercise generation functionalities, Lärka has been evolving into a research tool with a number of supportive modules for experimentation and visualization of research results, such as for selection of best corpus examples for language learners, for readability analysis of texts aimed at or produced by language learners, for prediction of single-word lexical difficulty, as well as for facilitating text-level annotation of language learner corpora, but also to collect data from exercises where learner interaction with the platform and their input have been used in research on metalinguistic awareness. Lärka is actively used in teaching grammar to university students, where we can report only those uses that we have explicitly been told about. As we do not require login to the platform, we do not know who our users are, but we can deduce from the logs that Lärka is being used beyond the reported schools and universities.

[1] https://spraakbanken.gu.se/larka

David Alfter, Lars Borin, Ildikó Pilán, Therese Lindström Tiedemann and Elena Volodina 2019. From Language Learning Platform to Infrastructure for Research on Language Learning. *Selected papers from the CLARIN Annual Conference 2018*. Linköping Electronic Conference Proceedings 159: 1–14.

Nowadays, Lärka is being adapted into a building block in an emerging second language research infrastructure SweLL (Volodina et al., 2018), within a larger context of the text-based research infrastructure developed by the national Swedish Language bank, Språkbanken, and SWE-CLARIN. This addresses an obvious need within CLARIN, as evidenced by the interest in the recent CLARIN workshop on "Interoperability of Second Language Resources and Tools".[2]

The current paper describes the new version of Lärka that was released in 2016, replacing the 2013 version, and illustrates improved and newly added functionalities.

2 Related work

There have been some attempts to combine exercise platforms with different types of data collection. The *Writing Mentor* Google Docs add-on, for example, allows users to get feedback on their writing in different categories such as coherence, topic development or use of sources to back up claims. The application uses natural language processing tools to provide users with feedback but at the same time collects the texts and all subsequent modifications to the texts that have been analyzed (Madnani et al., 2018). However, accessibility of the data for SLA research is limited.

The FeedBook project (Rudzewitz et al., 2017) is based on an English text book and presents the text book in a digitized interactive web platform that has been enriched with natural language processing to provide immediate fine-grained feedback to the students concerning both form and meaning errors. Teachers can also see their students' progress and provide individual feedback. The data is logged and is used iteratively for further improvement of the system, the data access so far being limited to the researchers involved in the project.

Most applications, however, are purely pedagogical. An outstanding example is the *Language Muse* Activity Palette (Burstein and Sabatini, 2016; Burstein et al., 2017). It allows teachers to upload texts and automatically generates exercises based on these texts. Texts are analyzed using natural language processing algorithms to identify different linguistic features such as multi-word expressions, syntactic relations and discourse structure. Based on the analysis, the platform creates over twenty different activities for the teacher to choose from, such as antonym exercises, homonym exercises or verb tense exercises. Teachers have full control over which texts are used, and are offered a possibility to edit automatically suggested exercise items. In that way, teachers can build a 'palette' of activities from the original text that best suits their and their students' needs.

Perez and Cuadros (2017) propose a framework for automatic exercise generation from user-specified texts that works with Spanish, Basque, English and French. Users can use texts of their own choosing in four different languages. The framework can generate three different kinds of tasks, namely gap exercises, multiple-choice exercises and sentence rearrangement exercises. Furthermore, the framework automatically generates hints for the gap exercise and allows for the adjustment of the number of distractors for multiple-choice exercises. Exercises are also exportable in Moodle's CLOZE[3] format, increasing its appeal.

On the other hand, there are multiple examples of SLA and psycholinguistic experiments that are staged through exercises that elicit certain types of data from language learners – data that helps researchers to address particular research questions, e.g. Andersson et al. (2018) investigating the influence of the native language on the processing of the word order in Swedish or Kerz and Wiechmann (2017) studying individual differences in L2 processing of multi-word phrases.

We argue that exercise generation platforms/applications have a capacity to mediate between language learners and researchers, bringing interests of the two groups together. We aim to foster this collaboration through the Lärka platform.

Lärka started as an exercise generation platform for learners of Swedish, and later it was extended to support the development and visualization of new algorithms in support of language learning. Now we are taking a new direction, combining research interests from Second Language Acquisition (SLA),

[2] See https://sweclarin.se/eng/workshop-interoperability-l2-resources-and-tools

[3] https://docs.moodle.org/23/en/Embedded_Answers_(Cloze)_question_type

Learner Corpus Research (LCR) and language learning into one and building an infrastructure supporting the collection of L2 data through exercises.

In the next sections, we delineate how Lärka can be and is used as a pedagogical tool in teaching students of Swedish linguistics (Sections 3.1 and 4), the different exercises in support of research aimed at L2 Swedish (Section 3.2), and the various components that constitute the research infrastructure facet of the platform (Section 5).

3 Lärka for learning and teaching

One of the main functionalities of Lärka is the automatic generation of exercises based on real-life authentic language examples from corpora. Exercise generation is aimed at two groups of learners: students of (Swedish) linguistics and learners of Swedish as a second or foreign language (L2).

3.1 Exercises for students of Linguistics

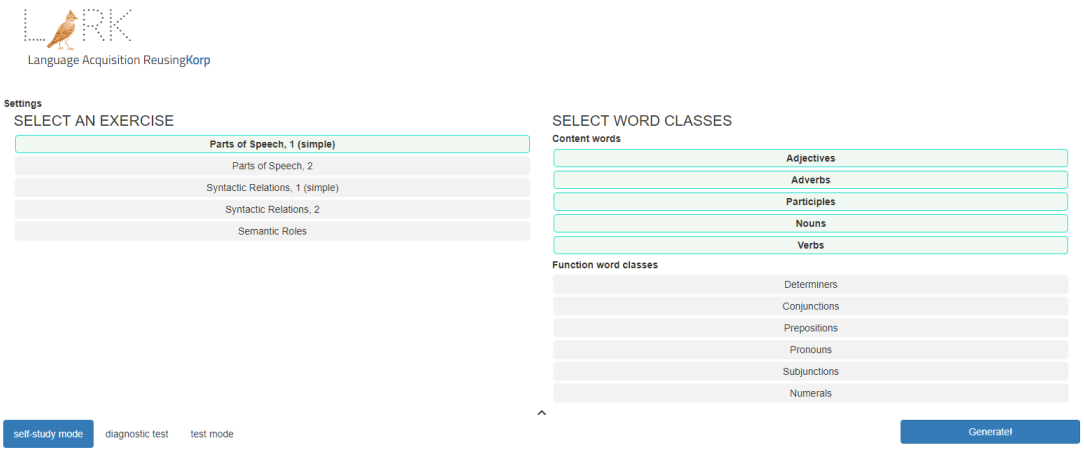

Figure 1: Exercises for linguists

Students learning grammatical analysis are in constant need of exercises and feedback on their analysis. Lärka offers exercises for linguistic analysis of parts of speech (word classes), syntactic relations and semantic roles. The exercises are based on authentic texts, which can make them more difficult than textbook examples. However, they are authentic examples of the type of texts the students are expected to be able to analyze in the future. Through on-the-spot feedback, students' learning is enhanced, especially if exercises are done at least partly in a class room setting with the possibility of consulting a teacher and/or the possibility of discussing one's analysis with a fellow student and together trying to make sense of why the automatic feedback said that they got it right or wrong (Lindström Tiedemann et al., 2016).

As mentioned above, Lärka offers students 3 types of exercises: parts of speech, syntactic relations and semantic roles. The first two offer two levels of difficulty (beginner and intermediate), whereas the third exercise, semantic roles, is only available as one level (Figure 1).

Exercises are presented as shown in Figure 2 with a sentence and a word or phrase highlighted in another colour. The learner then has to select from a multiple choice drop down box which answer is correct given the highlighted word or phrase. In part-of-speech exercises for example, learners have to select the correct part of speech for the highlighted word.

The exercises are available in three different modes: self-study, diagnostic test or test. Students can choose whichever mode they want to use. In self-study mode, answers can be revised as often as desired and needed, even after submitting the answer. In this way, if the answer was incorrect, it is possible to find the correct answer. In test mode, answers cannot be changed after submitting and the correct answer will be shown immediately after submitting. In diagnostic mode, as in test mode, answers cannot be changed after submitting. In addition, the number of exercise items is limited to three items of each main category, e.g. the part-of-speech exercise type covers eleven part-of-speech categories, resulting in a total of 33 diagnostic exercise items. Exercise generation stops after completion of all items in diagnostic

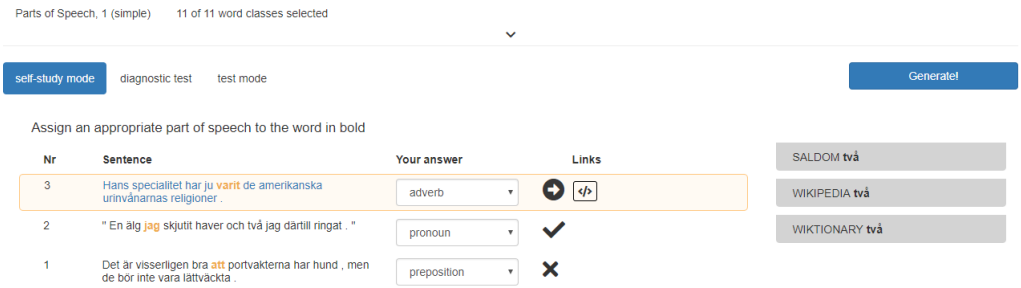

Figure 2: Execises for linguists

mode and a summary is provided which can be emailed to the teacher for further comments or to oneself in order to study the examples further or to be able to track one's learning. In contrast, the other two modes generate exercises infinitely. In both self-study and test mode the actual categories practiced can also be chosen (e.g. one can select to only practice adjectives and adverbs for part-of-speech exercises), whereas the diagnostic test automatically selects all available categories.

In order to avoid exercise item repetition, a sentence will be shown only once during the same session.

3.2 Exercises for language learners

Lärka offers a number of exercises for learners of L2 Swedish as illustrated in the following paragraphs. For all learner exercises, target vocabulary items are sampled from SVALex (François et al., 2016) and SweLLex (Volodina et al., 2016b). SVALex presents a list of lemmata occurring at the different CEFR (Common European Framework of Reference for Languages (Council of Europe, 2001)) levels in the textbook corpus COCTAILL (Volodina et al., 2014a). Similarly, SweLLex is based on the pilot SweLL corpus (Volodina et al., 2016a), a corpus of learner essays. We map each distribution to a single CEFR level according to two approaches, namely *first-occurrence* (Gala et al., 2013; Gala et al., 2014) and *threshold* (Alfter et al., 2016).

The exercises target lexical knowledge of Swedish L2 learners, and speaking pedagogically, train lexical knowledge from various points of view, namely: listening and spelling of lexical items, recognition of an appropriate item for a given context, morphological inflectional behaviour of individual lexical items, and linking definitions/translations with words. There are certainly a many other conceivable exercises that target different word knowledge aspects that we have not implemented. While even the exercise types that we currently offer are still in need of evaluation with teachers and learners, we do believe that they are useful. The session logs for the listening and spelling and word guess exercises show that there is interest in these types of exercises.

3.2.1 Vocabulary and inflection

Vocabulary exercises and inflection exercises have a multiple-choice format. Each item consists of a sentence containing a gap, as well as a list of five answer alternatives, of which one is correct and four are *distractors*, i.e. incorrect options (Figure 3). For vocabulary, distractors are chosen of the same word class as the target word. This morphological selection is further restricted by requiring that distractors be of the same number and/or definiteness as the target item for nouns or the same voice and/or tense for verbs. In case the restriction on the distractors returns too few results, these constraints can be relaxed or dropped.

For inflection exercises, we look up all morphological forms of the target word in Saldo's morphology (Borin et al., 2013a) and use a subset of those as distractors. Figures 3 and 4 show the vocabulary and inflection multiple choice exercise respectively.

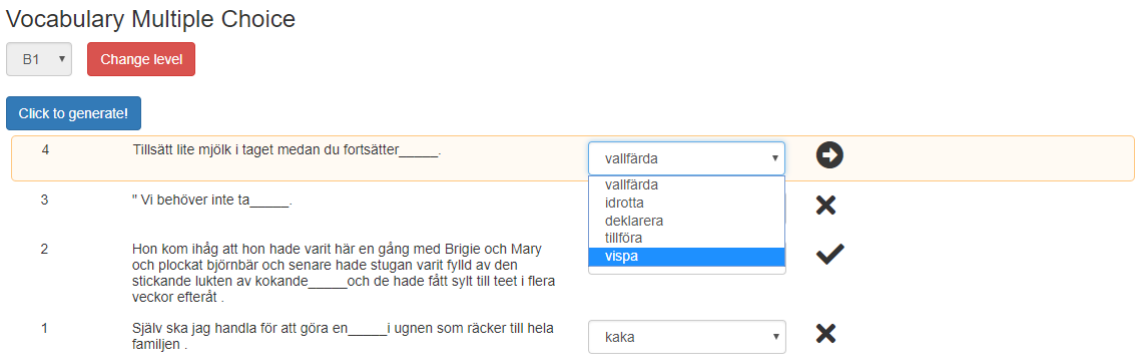

Figure 3: Vocabulary multiple choice

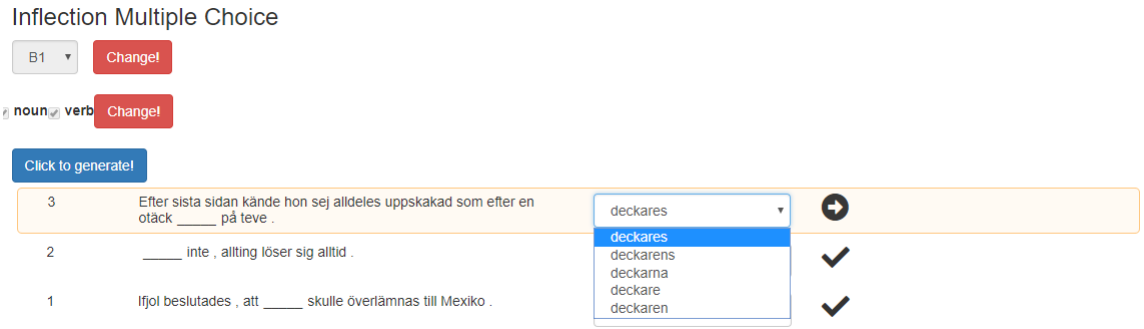

Figure 4: Inflection multiple choice

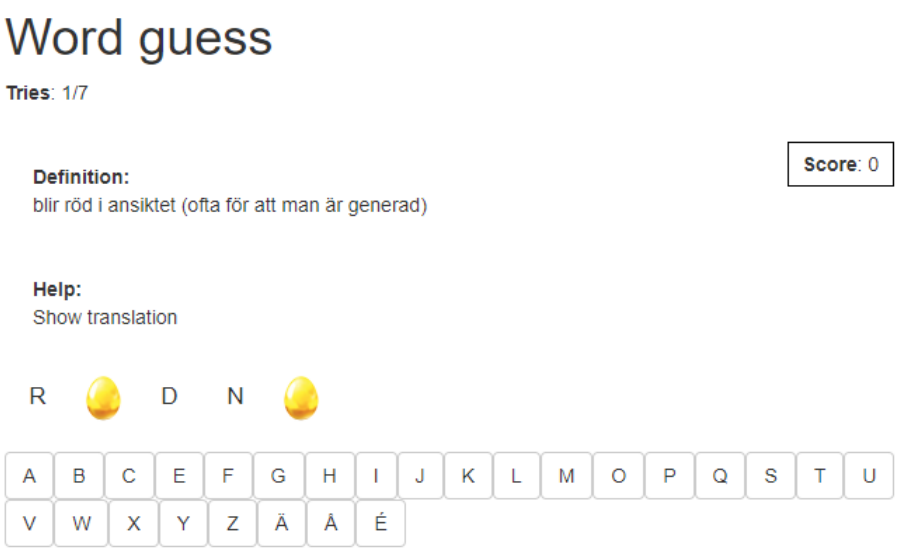

Figure 5: Word guess

3.2.2 Word guess

A recent addition to our platform is a simple word-level exercise, *Word guess*, that takes a step towards gamified learning. Word guess re-implements the well-known Hangman game format: users are presented with a number of hidden characters and the definition of the word in Swedish, and their task is to guess letters contained in the word, which eventually helps them guess the word itself, as shown in Figure 5.

Every time the guessed character is not in the word, users receive penalty points. In our learning-oriented version of the game, users can choose to receive clues such as the translation of the word (into a range of different languages). Both the definition and translations are retrieved from *Lexin*, a core-vocabulary lexicon for immigrants (Gellerstam, 1999). This game is a simple example of reusing information from lexical resources for gamified language learning activities.

3.2.3 Liwrix

Another exercise is the listening exercise *Liwrix* (Volodina and Pijetlovic, 2015). This exercise makes use of Text-to-Speech (TTS) technology by SitePal[4] to dynamically generate audio of single words and multi-word expressions. In the future, we also intend to include phrases and sentences, as was done in the previous version of Lärka. The delay is caused by the newly introduced hint system which needs to be modified in order to work with phrases and sentences.

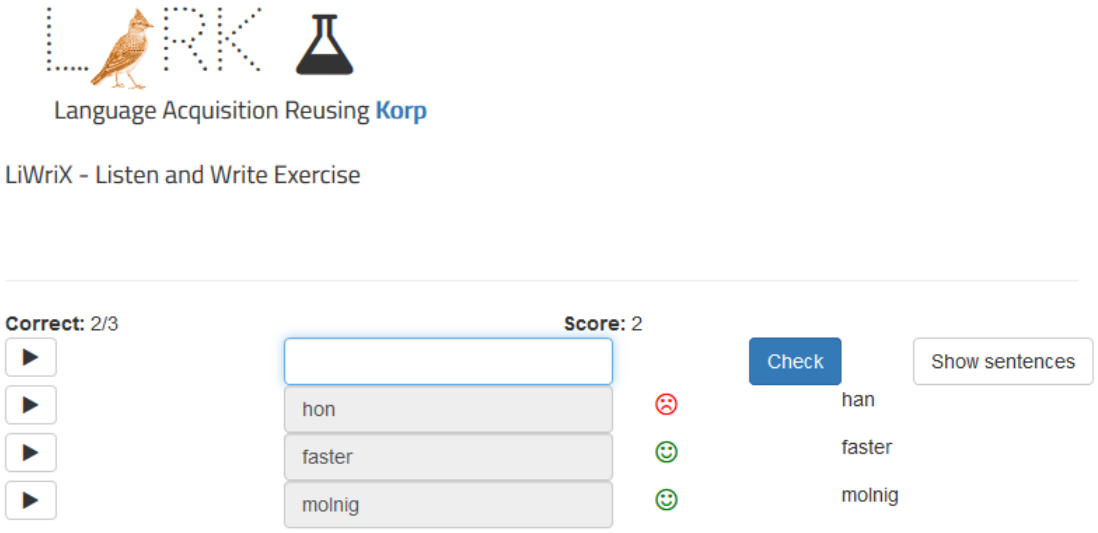

Figure 6: Liwrix

Figure 6 shows the exercise. By clicking on the button on the left, a word or multi-word expression is played and the answer is to be entered into the textfield. In addition, hints are available: As a first hint, but also to avoid problems with homonyms or possible mispronunciations, users can get “clues" in the form of a number of sentences in which the word(s) to be guessed appear in context. As a second hint, learners can choose to have the initial letter of the target word revealed.

Feedback is given in the form of a green smiley if the answer was correct and a red smiley if the answer was incorrect. In test mode (as in Figure 6) the correct answer is also shown irrespective of the correctness of the learner input.

4 Lärka in practice

Lärka for linguists has been used in introductions to grammar and linguistics in Sweden and Finland (Volodina et al., 2014b; Lindström Tiedemann et al., 2016). In Uppsala the platform was often used in lab sessions first so that students had a chance to consult a teacher when they had questions and they were also encouraged to discuss their analysis and the automatic feedback they got with their fellow students.

In Helsinki students have sometimes been encouraged to use it independently on courses in Swedish grammar where they have then been asked to hand in some of their analysis to their teacher or simply been told to use it to get more practice which is something they clearly cannot get too much of in learning grammatical analysis. Some exercise books might not even come with a key, which means that all exercises must be treated in class if the students are to find out what they did right or wrong.

[4] sitepal.com

In comparison, Lärka material is better suited in this case than many exercise books since it provides authentic texts accompanied by immediate automatic feedback.

The students felt that this was of great use and definitely thought that the platform should be used in the future. In a study with 45 students, Lärka was generally well received. Figure 7 shows that the majority of students were in favor of keeping Lärka as part of lab sessions with 34 students (78%) responding strongly in favor of keeping Lärka (scores 5–6), while 10 students (22%) showed more reservation (scores 3–4). No students voted against keeping Lärka (scores 1–2). Similarly, Figure 8 shows that 80% of students would recommend Lärka to a fellow student while 20% showed reservations.

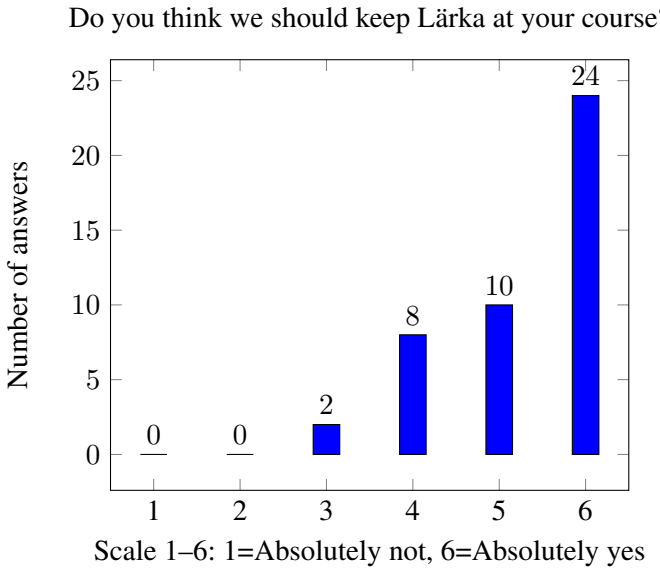

Figure 7: Evaluation results 1

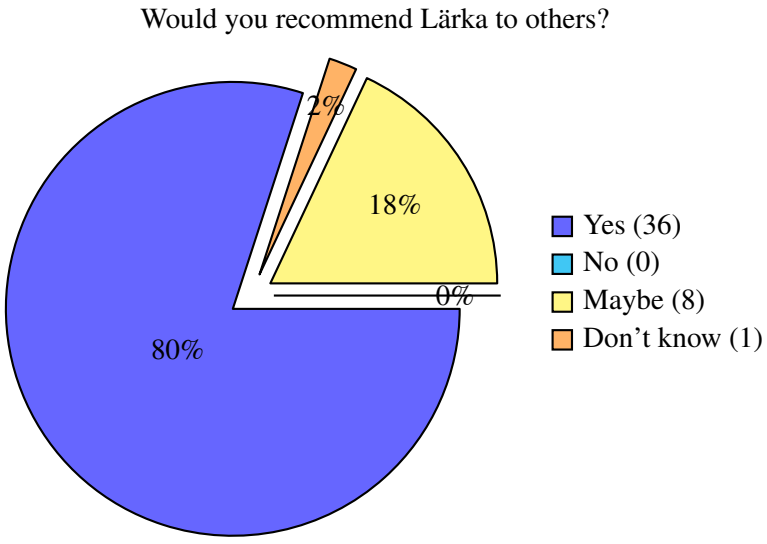

Figure 8: Evaluation results 2

A more recent analysis of the linguistic exercise log data collected through the 2016 version of the platform shows that during the time span from October 2016 to May 2018, there were 2086 sessions. One session is counted as a user using the platform from the moment of opening the page to closing it. As we do not require users to login, we create anonymous session identifiers each time a user opens the page. Thus, multiple sessions can stem from the same user. There were 126 sessions in the period from

(a)

Exercise type	# interactions
Part of Speech 1	28,544
Part of Speech 2	6,717
Semantic roles	553
Syntactic relations 1	8,426
Syntactic relations 2	2,842

(b)

Operating system	# interactions
Android	1,081
Linux	1,093
Mac OS X	11,516
Windows	18,962
iOs	2,102

Table 1: Interaction by exercise type (a) and operating system (b)

October 2016 to December 2016, 1449 sessions during 2017 and 511 sessions from January 2018 to May 2018.

During those 2086 sessions, a total of 47082 interactions were carried out. One interaction counts as an exercise item being completed. Interaction counts do not include self-corrections, mode changes or helps consulted. Table 1 (a) shows the breakdown of the interactions per exercise type. A logging feature that was added later[5] was the logging of whether the page was accessed from a mobile device and which operating system was used to access the page. The logs show that the linguist exercise was accessed 3,184 times (~10%) from a mobile device, as opposed to 31,571 times from a non-mobile device. Table 1 (b) shows the breakdown of interactions by operating system.

Furthermore, we can see that the platform was mainly accessed from Sweden (91%) and Finland (8%), but also from other countries such as the US, Poland, Germany, the Netherlands, Turkey, Estonia, the UK, India, Belgium, Switzerland, Japan, Canada and Russia, together making up the remaining 1%.

5 Lärka as research infrastructure

Lärka is being developed to serve as one of the e-infrastructure components offered to the research community by the Swedish CLARIN B-centre Språkbanken Text at the University of Gothenburg. Specifically Lärka is intended to be used as an infrastructure for research in (Swedish as) L2 acquisition. Currently Lärka offers modules for (1) collection of data from learners through their interaction with the platform, i.e. exercise logs; (2) text-level annotation of learner essays and course book texts; as well as (3) experimentation and visualization of the ongoing research in support of language learning.

With these modules, materials and exercises can be tailored drawing on vast collections of naturally occurring language, in a precise yet flexible as well as replicable way, and students' responses and reactions can be recorded in detail for subsequent quantitative and qualitative analysis. In order to achieve the necessary combination of precision and flexibility, we integrate natural language processing tools and algorithms for corpus example selection, text assessment and automatic exercise generation. These aspects are described in more detail below. A recent direction is "profiling" lexical and grammatical competences that learners of Swedish have, where we experiment with different lexical resources for exercise creation, and in the near future expect to integrate research on grammar profiles.[6]

5.1 Corpus example selection

In Lärka, the automatically generated exercises for language learners rely on *HitEx* (*Hitta Exempel* 'find examples'), a tool for selecting and ranking corpus examples (Pilán et al., 2017). The main purpose of HitEx is to identify sentences from generic corpora which are suitable as exercise items for L2 learners. The suitability of the sentences is determined based on a number of parameters that reflect different linguistic characteristics of the sentences. Through a graphical user interface, it is also possible to conduct a sentence search based on parameters customized by the user. The selection criteria include a wide variety of linguistic aspects such as the desired difficulty level based on CEFR, typicality based on word

[5] That is why the total is lower than 47,082

[6] `https://spraakbanken.gu.se/eng/l2-profiling`

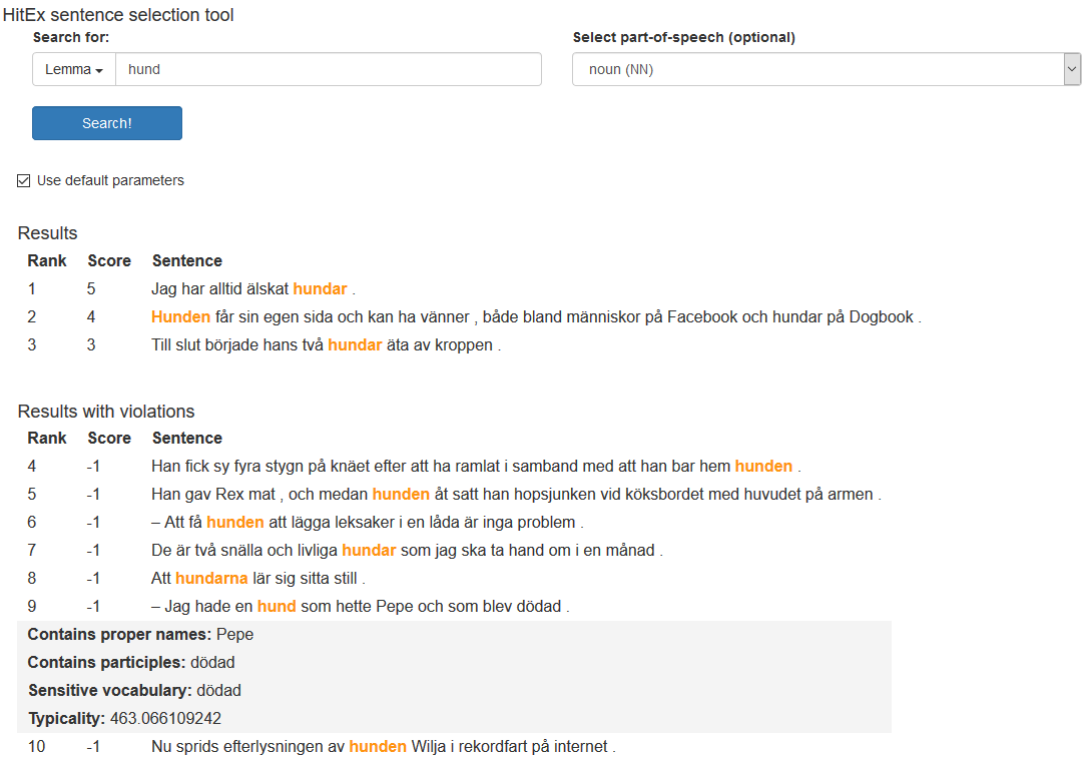

Figure 9: Corpus example selection tool HitEx: Results

co-occurrence measures, as well as the absence of anaphoric expressions and sensitive vocabulary (e.g. profanities), just to name a few. It is also possible to use a set of default parameters for searching. Figure 9 shows the results of HitEx. Sentences which fulfill all the required parameter constraints are shown on top while results that violate one or more constraints are shown under 'Results with violations'. Upon clicking on one of the sentences, more information is shown.

5.2 Text complexity evaluation

Another functionality, *TextEval*, offers an interface to automatically assess Swedish texts for their degree of complexity according to the CEFR. Texts can be either learner productions (e.g. essays) or texts written by experts as reading material for learners. The machine learning based automatic analysis returns an overall CEFR level for the text, as well as a list of linguistic indicators relevant for measuring text complexity, such as the average length of sentences and tokens, LIX score and nominal ratio. In addition, it is possible to add a color-enhanced highlighting for words per CEFR levels which provides users with a straightforward visual feedback about the lexical complexity of a text. Figure 10 shows the analysis of a text with word-level CEFR highlighting. We use the aforementioned lists SVALex and SweLLex to mark up receptive and productive vocabulary respectively. For each CEFR level, a darker and a lighter shade of the same color represents productive and receptive vocabulary respectively at the given level.

5.3 Lexical complexity prediction

Based on the word lists SVALex and SweLLex, which have been transformed so as to map each word to a single CEFR level as described in Alfter et al. (2016), we have built a module capable of predicting the complexity of any Swedish word, not only words occurring in the word lists (Alfter and Volodina, 2018). For each word, we extract both traditional word-based features such as length, number of syllables, number of homonyms and also information about topics, i.e. which topics a word belongs to. For example, the word *fisk* 'fish' would occur in the topics 'Animals' and 'Food'. We then feed a machine

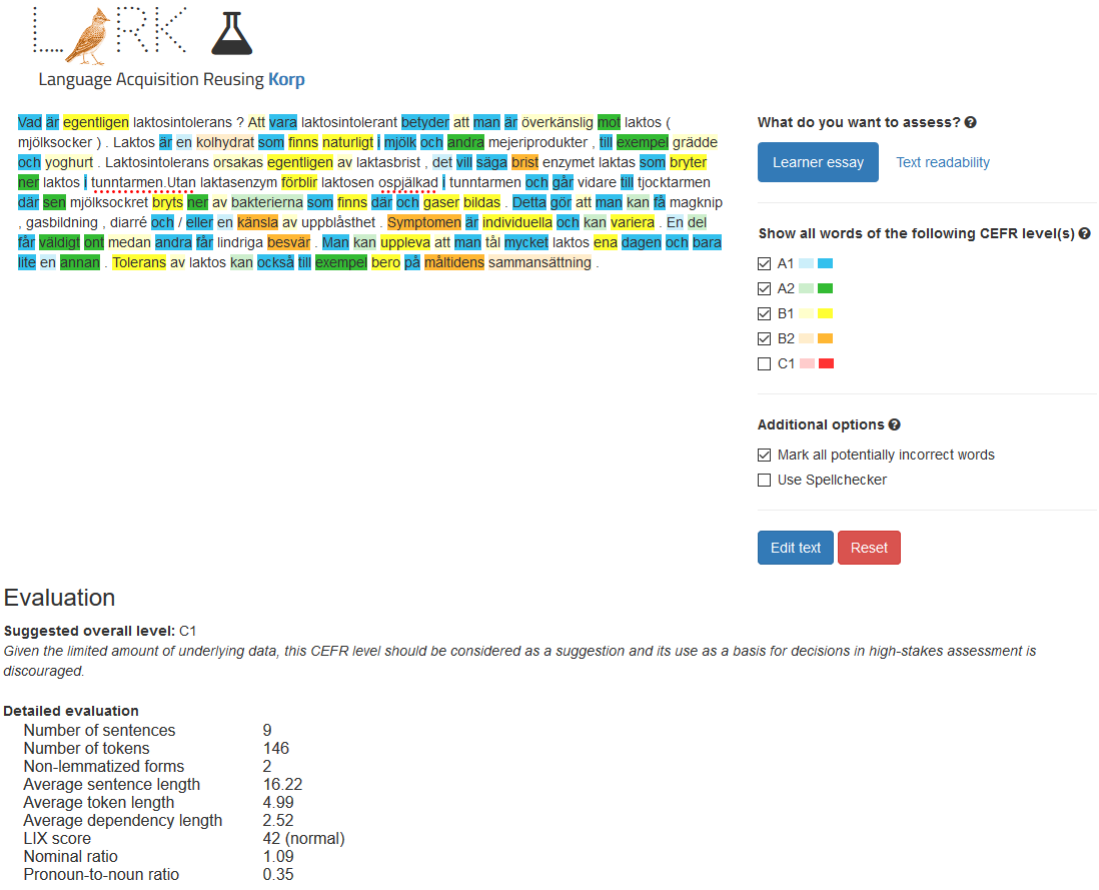

Figure 10: Text complexity evaluation tool TextEval

learning algorithm these feature vectors as well as the predicted mapped single CEFR level of the word and let the algorithm learn how to map from these features to CEFR levels.

An interested user can test a bespoke interface to get predictions about the complexity of a word and its target level (receptive versus productive), as shown in Figure 11. This user interface can be used for getting predictions of any word, not only words present in the word lists. The input word is transformed into a feature vector as described above and then fed into the classifier, which predicts a label. Figure 11 shows the predictions for *hund* 'dog', *vovve* 'doggy' (childish or endearing term for 'dog') and *byracka* 'mutt' (derogatory term for 'dog').

5.4 Annotation editor

Lärka contains an annotation editor that can be used for XML markup of textbooks. The editor provides an intuitive menu that makes adding XML tags easy. The editor keeps track of current settings in order to make adding new elements as easy as possible. It also automatically increments lesson counters and other counters. The editor offers the possibility to download the annotated text as an XML file. The current version of the editor also includes the possibility to save one's progress and continue working on it at a later moment in time without the need to login. The SweLL corpus pilot project (Volodina et al., 2016a) and the COCTAILL corpus project (Volodina et al., 2014a) used a previous version of the annotation editor to achieve consistent XML markup of essays and course books as well as to simplify the annotation process by providing an intuitive and intelligent user interface.

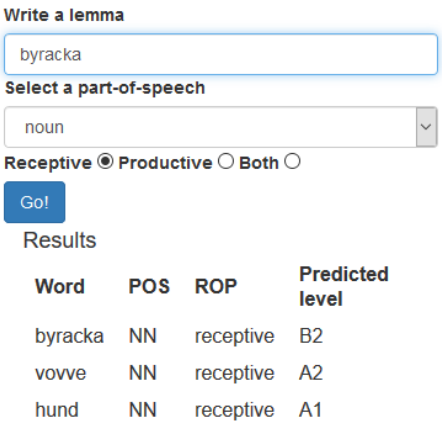

Figure 11: User interface for lexical complexity prediction

5.5 Lexicographic annotation tool

Another annotation tool that has recently been added to Lärka is the Lexicographic Annotation Tool, Legato. This tool can be used to annotate words or word senses on different lexicographic levels. Figure 12 shows the tool in the 'register' annotation mode. Here, the annotator is presented with a SALDO sense (viz. *gammal* 'old'), its part-of-speech (adjective) and the predicted CEFR level (A1). In addition, the tool shows the primary and secondary SALDO descriptors, if available. As different senses of a word can still be ambiguous as to the category to be annotated, we also show an example sentence where the word sense is highlighted, in this case surrounded by two asterisks (**). The example sentences have been selected to be of the same CEFR level as the word sense in question.

The main part of the interface shows the annotation possibilities. In the example shown, different options for register are shown. The annotator can select none, one, or more than one of these possibilities.

Finally, using the buttons at the bottom, annotators can leave the interface to annotate either another lexicographic category or to stop annotating altogether. Items can be skipped if the annotator is unsure about the annotation. In this case, the item will be added to the list of skipped items which can be accessed by clicking the button on top next to the 'Guidelines' button. This opens up a side menu which shows all the skipped items. By clicking on any of these items, the interface returns to the item in question. The interface also offers a search functionality which makes searching through the list of items easy.

In addition, the interface keeps track of different annotators and their progress across different annotation categories. Thus, if an annotator annotates ten items in 'morphology', then returns to the main screen and annotates ten items in 'nominal gender', then returns to morphology, the interface will resume at item number eleven. This also works across sessions. Thus, annotation does not have to be done in one fell swoop but can be done intermittently. The skipped items are also saved per annotator and category. For example, if annotator A skips *gammal* 'old' in 'register' but not in 'morphology', it will turn up for annotator A under 'register' until it is resolved. All data is saved to a data base on the server.

Besides fully manual annotation, the tool also offers a semi-automatic annotation mode where some of the values have been automatically extracted by linking together various resources. In this annotation mode, if values have been found, the annotator's task is to check whether the values are correct and correct them if necessary. If no values have been found, the annotator proceeds as in manual mode.

6 Ongoing work and planned extensions

Besides the activities described in this paper, the addition of new exercise formats and the implementation of a diagnostic placement test are currently under development. In the near future we plan to add a login functionality as well as an infrastructure to log more specific user data. This would enable us to create a

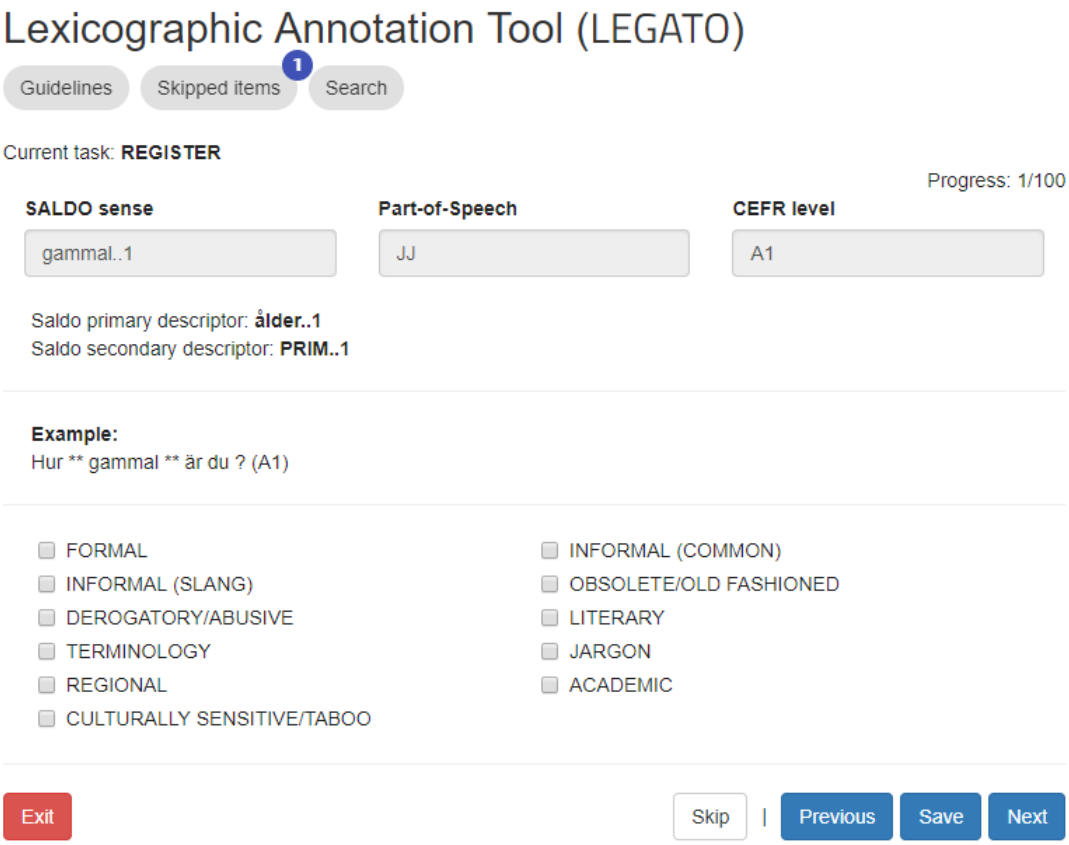

Figure 12: Lexicographic annotation tool Legato

valuable resource for modeling learners (e.g. L1-specific errors, learners' development over time) and to offer adaptive exercises.

References

David Alfter and Elena Volodina. 2018. Towards single word lexical complexity prediction. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 79–88.

David Alfter, Yuri Bizzoni, Anders Agebjörn, Elena Volodina, and Ildikó Pilán. 2016. From distributions to labels: A lexical proficiency analysis using learner corpora. In *Proceedings of the joint workshop on NLP for Computer Assisted Language Learning and NLP for Language Acquisition at SLTC, Umeå, 16th November 2016*, number 130, pages 1–7. Linköping University Electronic Press.

Annika Andersson, Susan Sayehli, and Marianne Gullberg. 2018. Language background affects online word order processing in a second language but not offline. *Bilingualism: Language and Cognition*, pages 1–24.

Lars Borin and Anju Saxena. 2004. Grammar, incorporated. In Peter Juel Henrichsen, editor, *CALL for the Nordic languages*, pages 125–145. Samfundslitteratur, Copenhagen.

Lars Borin, Markus Forsberg, and Johan Roxendal. 2012. Korp-the corpus infrastructure of Spräkbanken. In *PLRroceedings of EC*, pages 47 20124–478.

Lars Borin, Markus Forsberg, and Lennart Lönngren. 2013a. SALDO: a touch of yin to WordNets yang. *Language Resources and Evaluation*, 47(4):1191–1211.

Lars Borin, Markus Forsberg, Leif-Jöran Olsson, Olof Olsson, and Jonatan Uppström. 2013b. The lexical editing system of Karp. In *Proceedings of the eLex 2013 conference*, pages 503–516.

Lars Borin, Nina Tahmasebi, Elena Volodina, Stefan Ekman, Caspar Jordan, Jon Viklund, Beáta Megyesi, Jesper Näsman, Anne Palmér, Mats Wirén, Kristina Nilsson Björkenstam, Gintar Grigonyt, Sofia Gustafson Capková, and Tomasz Kosiski. 2017. Swe-Clarin: Language resources and technology for digital humanities. *Digital Humanities 2016. Extended Papers of the International Symposium on Digital Humanities (DH 2016) Växjö, Sweden, November, 7-8, 2016. Edited by Koraljka Golub, Marcelo Milra*, Vol-2021.

Jill Burstein and John Sabatini. 2016. The language muse activity palette. *Adaptive educational technologies for literacy instruction*, pages 275–280.

Jill Burstein, Nitin Madnani, John Sabatini, Dan McCaffrey, Kietha Biggers, and Kelsey Dreier. 2017. Generating language activities in real-time for English learners using language muse. In *Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale*, pages 213–215. ACM.

Council of Europe. 2001. *Common European Framework of Reference for Languages: Learning, Teaching, Assessment*. Press Syndicate of the University of Cambridge.

Thomas François, Elena Volodina, Ildikó Pilán, and Anaïs Tack. 2016. SVALex: a CEFR-graded lexical resource for Swedish foreign and second language learners. In *LREC*.

Núria Gala, Thomas François, and Cédrick Fairon. 2013. Towards a French lexicon with difficulty measures: NLP helping to bridge the gap between traditional dictionaries and specialized lexicons. *E-lexicography in the 21st century: thinking outside the paper., Tallin, Estonia*.

Núria Gala, Thomas François, Delphine Bernhard, and Cédrick Fairon. 2014. Un modèle pour prédire la complexité lexicale et graduer les mots. In *TALN 2014*, pages 91–102.

Martin Gellerstam. 1999. LEXIN-lexikon för invandrare. *LexicoNordica*, (6).

Elma Kerz and Daniel Wiechmann. 2017. Individual differences in l2 processing of multi-word phrases: Effects of working memory and personality. In *International Conference on Computational and Corpus-Based Phraseology*, pages 306–321. Springer.

Therese Lindström Tiedemann, Elena Volodina, and Håkan Jansson. 2016. Lärka: ett verktyg för träning av språkterminologi och grammatik. *LexicoNordica*, 23:161–181.

Nitin Madnani, Jill Burstein, Norbert Elliot, Beata Beigman Klebanov, Diane Napolitano, Slava Andreyev, and Maxwell Schwartz. 2018. Writing mentor: Self-regulated writing feedback for struggling writers. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 113–117.

Naiara Perez and Montse Cuadros. 2017. Multilingual call framework for automatic language exercise generation from free text. In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 49–52.

Ildikó Pilán, Elena Volodina, and Lars Borin. 2017. Candidate sentence selection for language learning exercises: from a comprehensive framework to an empirical evaluation. *TAL*, 57(3/2016):67–91.

Björn Rudzewitz, Ramon Ziai, Kordula De Kuthy, and Detmar Meurers. 2017. Developing a web-based workbook for english supporting the interaction of students and teachers. In *Proceedings of the Joint 6th Workshop on NLP for Computer Assisted Language Learning and 2nd Workshop on NLP for Research on Language Acquisition at NoDaLiDa, Gothenburg, 22nd May 2017*, number 134, pages 36–46, Linköping. Linköping University Electronic Press.

Elena Volodina and Dijana Pijetlovic. 2015. Lark trills for language drills: Text-to-speech technology for language learners. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 107–117.

Elena Volodina, Lars Borin, Hrafn Lofsson, Birna Arnbjörnsdóttir, and Guðmundur Örn Leifsson. 2012. Waste not; want not: Towards a system architecture for ICALL based on NLP component re-use. In *Proceedings of the SLTC 2012 workshop on NLP for CALL; Lund; 25th October; 2012*, number 080, pages 47–58. Linköping University Electronic Press.

Elena Volodina, Ildikó Pilán, Stian Rødven Eide, and Hannes Heidarsson. 2014a. You get what you annotate: a pedagogically annotated corpus of coursebooks for Swedish as a second language. In *Proceedings of the third workshop on NLP for computer-assisted language learning at SLTC 2014, Uppsala University*, number 107. Linköping University Electronic Press.

Elena Volodina, Ildikó Pilán, Lars Borin, and Therese Tiedemann Lindström. 2014b. A flexible language learning platform based on language resources and web services. In *Proceedings of LREC 2014, Reykjavik, Iceland*, pages 3973–3978.

Elena Volodina, Ildikó Pilán, Ingegerd Enström, Lorena Llozhi, Peter Lundkvist, Gunlög Sundberg, and Monica Sandell. 2016a. SweLL on the rise: Swedish learner language corpus for european reference level studies. *arXiv preprint arXiv:1604.06583*.

Elena Volodina, Ildikó Pilán, Lorena Llozhi, Baptiste Degryse, and Thomas François. 2016b. SweLLex: second language learners productive vocabulary. In *Proceedings of the joint workshop on NLP for Computer Assisted Language Learning and NLP for Language Acquisition at SLTC, Umeå, 16th November 2016*, number 130, pages 76–84. Linköping University Electronic Press.

Elena Volodina, Lena Granstedt, Sofia Johansson, Beáta Megyesi, Julia Prentice, Dan Rosén, Carl-Johan Schenström, Gunlög Sundberg, and Mats Wirén. 2018. Annotation of learner corpora: first swell insights. *Proceedings of SLTC 2018*.