

Natural Language Processing and Language Learning

Detmar Meurers
Universität Tübingen
dm@sfs.uni-tuebingen.de

As a relatively young field of research and development started by work on cryptanalysis and machine translation around 50 years ago, Natural Language Processing (NLP) is concerned with the automated processing of human language. It addresses the analysis and generation of written and spoken language, though speech processing is often regarded as a separate subfield. NLP emphasizes processing and applications and as such can be seen as the applied side of Computational Linguistics, the interdisciplinary field of research concerned with formal analysis and modeling of language and its applications at the intersection of Linguistics, Computer Science, and Psychology. In terms of the language aspects dealt with in NLP, traditionally lexical, morphological and syntactic aspects of language were at the center of attention, but aspects of meaning, discourse, and the relation to the extra-linguistic context have become increasingly prominent in the last decade. A good introduction and overview of the field is provided in Jurafsky & Martin (2009).

This article explores the relevance and uses of NLP in the context of language learning, focusing on written language. For a recent overview of technology targeting pronunciation, see Pennington & Rogerson-Revell (2019, ch. 5). We will focus on motivating the relevance, characterizing the techniques, and delineating the uses of NLP. More historical background and discussion can be found in Nerbonne (2003), Heift & Schulze (2007, 2015), and Heift (2017).

We can distinguish two broad uses of NLP related to language learning: On the one hand, NLP can be used to *analyze learner language*, i.e., words, sentences, or texts produced by language learners. This includes the development of NLP techniques for the analysis of learner language by tutoring systems in Intelligent Computer-Assisted Language Learning (ICALL, cf. Heift, 2017), automated scoring in language testing, as well as the analysis and annotation of learner corpora (cf. Granger, this volume).

On the other hand, NLP for the *analysis of native language* can also play an important role in the language learning context. Applications in this second domain support the search for and the enhanced presentation of native language reading material for language learners as well as the generation of exercises and tests based on authentic materials.

1 NLP and the Analysis of Learner Language

Intelligent Language Tutoring Systems (ILTS) use NLP to provide individualized feedback to learners working on activities, usually in the form of workbook-style exercises as in the E-Tutor (Heift, 2010), Robo-Sensei (Nagata, 2009), TAGARELA (Amaral & Meurers, 2011), the i-tutor (Choi, 2016), and the FeedBook (Meurers et al., 2019). The NLP analysis may also be used to individually adjust the sequencing of the material and to update the

learner model (cf. Schulze, 2011). Typically the focus of the analysis is on form errors made by the learner, even though in principle feedback can also highlight correctly used forms or target aspects of meaning or the appropriateness of a learner response given the input provided by an exercise.

What motivates the use of NLP in a tutoring system? To be able to provide feedback and keep track of abilities in a learner model, an ILTS must obtain information about the student's abilities. How this can be done depends directly on the nature of the activity and the learner responses they support, i.e., the ways in which they require the learner to produce language or interact with the system. The interaction between the activity or task given to a learner and the learner response is an important topic in language assessment (cf. Bachman & Palmer, 1996) and task-based language teaching and learning (Ellis, 2009) and it arguably is crucial for determining the system analysis requirements of different activity types (Quixal & Meurers, 2016).

For exercises that explicitly or implicitly require the learner to provide responses using forms from a small, predefined set, it often is possible to anticipate all potential well-formed and ill-formed learner responses, or at least the most common ones given a particular learner population. The intended system feedback for each case can then be explicitly specified for each potential response. In such a setup, knowledge about language and learners is exclusively expressed in this extensional mapping provided offline when the exercise is created. Sometimes targets are allowed to include regular expressions (http://en.wikipedia.org/wiki/Regular_expression) to support a more compact specification. The online process of comparing an actual learner response to the anticipated targets is a simple string comparison requiring no linguistic knowledge and thus no NLP. Correspondingly, the quiz options of general Course Management Systems (Moodle, ILIAS, etc.) can support such language exercises just as they support quizzes for math, geography, or other subjects. The same is true of the general web tools underlying the many quiz pages on the web (e.g., <http://www.eslcafe.com/quiz/>). Tools such as Hot Potatoes (<http://hotpot.uvic.ca/>) make it easier to specify some exercise forms commonly used in language teaching, but also use the same general processing setup without NLP.

For many types of language learning activities, however, extensionally specifying a direct and complete mapping between potential learner input and intended feedback is not feasible. Nagata (2009, pp. 563f) provides a clear illustration of this with an exercise taken from her Japanese tutor ROBO-SENSEI in which the learner reads a short communicative context and is asked to produce a sentence in Japanese that is provided in English by the system. The learner response in this exercise is directly dependent on the input provided by the exercise (a *direct response* in the terminology of Bachman & Palmer, 1996), so that a short, seven word sentence can be defined as target answer. Yet after considering possible well-formed lexical, orthographic, and word order variants, one already obtains 6048 correct sentences which could be entered by the learner. Considering incorrect options, even if one restricts ill-formed patterns to wrong particle and conjugation choices, one obtains almost a million sentences. Explicitly specifying a mapping between a million anticipated responses and their corresponding feedback clearly is infeasible. Note that the explosion of possible learner answers illustrated by Nagata already is a problem for the direct response in a constrained activity, where the meaning to be expressed was fixed and only form variation was anticipated. A further dimension of potential variation in responses arises when going beyond the analysis of language as a system (parallel to *system-referenced* tests in language assessment, cf. Baker, 1989) to an analysis of the ability to use language to appropriately complete a given task (*performance-referenced*).

In conclusion, for the wide range of language activities supporting significant well-formed or ill-formed variation of form, meaning, or task-appropriateness of learner responses, it is necessary to abstract away from the specific string entered by the learner to more general classes of properties by automatically analyzing the learner input using NLP algorithms and resources. Generation of feedback, learner model updating, and instructional sequencing can then be based on the small number of language properties and categories derived through NLP analysis instead of on the large number of string instances they denote.

What is the nature of the learner language properties to be identified? Research on Intelligent Language Tutors has traditionally focused on *learner errors*, identifying and providing feedback on them. Automatic analysis can also be used in an ILTS to identify well-formed language properties to be able to provide positive feedback or record in a learner model that a given response provided evidence for the correct realization of a particular construction, lexical usage, or syntactic relation. All approaches to detecting and diagnosing learner errors must explicitly or implicitly model the space of well-formed and ill-formed variation that is possible given a particular activity and a given learner. Insights into activity design and the language development of learners thus are crucial for effective NLP analysis of learner errors.

Errors are also present in native language texts and the need to develop *robust NLP*, which also works in suboptimal conditions (due to unknown or unexpected forms and patterns, noise), has been a driving force behind the shift from theory-driven, rule-based NLP in the 80s and 90s to the now-dominant data-driven, statistical and machine learning approaches. However, there is an important difference in the goal of the NLP use in an ILTS compared to that in other NLP domains. NLP is made robust to gloss over errors and unexpected aspects of the system input with the goal of producing some result, such as a syntactic analysis returned by a parser, or a translation provided by a machine translation system. The traditional goal of the NLP in an ILTS, on the other hand, is to identify the characteristics of learner language and in which way the learner responses diverge from the expected targets in order to provide feedback to the learner. So errors here are the goal of the abstraction performed by the NLP, not something to be glossed over by robustness of processing.

Writer's aids such as the standard spell and grammar checkers (Dickinson, 2006) share the ILTS focus on identifying errors, but they rely on assumptions about typical errors made by native speakers which do not carry over to language learners. For example, Rimrott & Heift (2008) observe that "in contrast to most misspellings by native writers, many L2 misspellings are multiple-edit errors and are thus not corrected by a spell checker designed for native writers." Tschichold (1999) also points out that traditional writer's aids are not necessarily helpful for language learners since learners need more scaffolding than a list of alternatives from which to choose. Writer's aids tools targeting language learners, such as the ESL Assistant (Gamon et al., 2009), therefore provide more feedback and, e.g., concordance views of alternatives to support the language learner in understanding the alternatives and choosing the right. The goal of writer's aids is to support the second language user in writing a functional, well-formed text, not to support them in acquiring the language as is the goal of an ILTS. Where writing well-formed and well-structured texts is the goal, advanced learners can also benefit from the quickly developing market for automatic writing evaluation tools, such as <https://writeandimprove.com>, <http://noredink.com>, <http://grammarly.com>, or <http://criterion.ets.org>.

NLP methods for the diagnosis of learner errors fall into two general classes: On the one hand, most of the traditional development has gone into ***language licensing approaches*** which analyze the entire learner response. On the other hand, there is a growing number of

pattern-matching approaches which target specific error patterns and types (e.g., preposition or determiner errors) ignoring any learner response or part thereof that does not fit the pattern.

Language licensing approaches are based on formal grammars of the language to be licensed, which can be expressed in one of two general ways (cf. Johnson, 1994). In a *validity-based* setup, a grammar is a set of rules and recognizing a string amounts to finding valid derivations. Simply put, the more rules are added to the grammar, the more types of strings can be licensed; if there are no rules, nothing can be licensed. In a *satisfiability-based* setup, a grammar is a set of constraints and a string is licensed if its model satisfies all constraints in the grammar. Thus the more constraints are added, the fewer types of strings are licensed; if there are no constraints in the grammar, any string is licensed.

Corresponding to these two types of formal grammars, there essentially are two types of approaches to analyzing a string with the goal of diagnosing learner errors. The *mal-rule approach* follows the validity-based perspective and uses standard parsing algorithms. Starting with a standard native language grammar, rules are added to license strings which are used by language learners but not in the native language, i.e., so-called *mal-rules* used to license learner errors (cf., e.g., Sleeman, 1982; Matthews, 1992, and references therein). Given that a specific error type can manifest itself in a large number of rules – e.g., an error in subject-verb agreement can appear in all rules realizing subjects together with a finite verb – *meta-rules* can be used to capture generalizations over rules (Weischedel & Sondheimer, 1983). The *mal-rule* approach can work well when errors correspond to the local tree licensed by a single grammar rule. Otherwise, the interaction of multiple rules must be taken into account, which makes it significantly more difficult to identify an error and to control the interaction of *mal-rules* with regular rules. To reduce the search space resulting from rule interaction, the use of *mal-rules* can be limited. In the simple case, the *mal-rules* are only added after an analysis using the regular grammar fails. Yet this only reduces the search space for parsing well-formed strings; if parsing fails, the question remains which *mal-rules* need to be added. The ICICLE system (Michaud & McCoy, 2004) presents an interesting solution by selecting the groups of rules to be used based on *learner modeling*. It parses using the native rule set for all structures which the learner has shown mastery of. For structures assumed to currently be acquired by the learner, both the native and the *mal-rules* are used. And for structures beyond the developmental level of the learner, neither regular nor *mal-rules* are included. When moving from traditional parsing to parsing with probabilistic grammars, one obtains a further option for distinguishing native from learner structures by inspecting the probabilities associated with the rules (cf. Wagner & Foster, 2009, and references therein). Different from such statistical approaches based on rules and the insights they capture, the most recent research on grammatical error correction (GEC) mostly focuses on detecting and correcting errors in written text as a type of translation problem. The idea is to map ill-formed to well-formed text directly using statistical machine translation methods (Junczys-Dowmunt & Grundkiewicz, 2016) or current neural network approaches (Chollampatt & Ng, 2018), which is the quantitative state of the art for GEC, but very limited in relevance for research and applications for which linguistic rules and (mis)conceptualization play a role.

The second group of language licensing approaches is typically referred to as *constraint relaxation* (Kwasny & Sondheimer, 1981), which is an option in a satisfiability-based grammar setup or when using rule-based grammars with complex categories related through unification or other constraints which can be relaxed. When parsing is treated as a general constraint satisfaction problem, general purpose conflict detection algorithm can be used to diagnose learner errors (Boyd, 2012). The idea of constraint relaxation is to eliminate certain constraints from the grammar, e.g., specifications ensuring agreement, thereby allowing the

grammar to license more strings than before. This assumes that an error can be mapped to a particular constraint to be relaxed, i.e., the domain of the learner error and that of the constraint in the grammar must correspond closely. Instead of eliminating constraints outright, constraints can also be associated with weights controlling the likelihood of an analysis (Foth et al., 2005), which raises the interesting issue how such flexible control can be informed by the ranking of errors likely to occur for a particular learner given a particular task. Other proposals combine constraint relaxation with aspects of *mal*-rules. Reuer (2003) combines a constraint relaxation technique with a standard parsing algorithm modified to license strings in which words have been inserted or omitted, an idea which essentially moves generalizations over rules in the spirit of meta-rules into the parsing algorithm.

For **pattern-matching**, the most common approach is to match a typical **error pattern**, such as the pattern looking for *of cause* in place of *of course*. By performing the pattern matching on the part-of-speech tagged, chunked, and sentence delimited learner string, one can also specify error patterns such as that of a singular noun immediately preceding a plural finite verb (e.g., in *The baseball team are established.*). This approach is commonly used in standard grammar checkers and e.g., realized in the open source LanguageTool (Naber, 2003; <http://www.languagetool.org>), which was not developed with language learners in mind, but readily supports the specification of error patterns that are typical for particular learner populations.

Alternatively, pattern-matching can also be used to identify **contexts patterns** that are likely to include errors. For example, determiners are a well-known problem area for certain learners of English. Using a pattern which identifies all nouns (or all noun chunks) in the learner response, one can then make a prediction about the correct determiner to use for this noun in its context and compare this prediction to the determiner use (or lack thereof) in the learner response. Given that determiners and preposition errors are among the most common English learner errors found and that the task lends itself well to the current machine learning approaches in computational linguistics (and raises the interesting general question how much context and which linguistic generalizations are needed for predicting such functional elements), these error types have received particular attention (cf., e.g., De Felice, 2008, and references therein).

Complementing the analysis of form, for an ILTS to offer meaning-based, contextualized activities it is important to provide an **automatic analysis of meaning** aspects, e.g., to determine whether the answer given by the learner for a reading comprehension question makes sense given the reading. While most NLP work in ILTS has addressed form issues, some work has addressed the analysis of meaning (Delmonte, 2003; Ramsay & Mirzaiean, 2005; Bailey & Meurers, 2008) and the issue is directly related to work in computer-assisted assessment systems outside of language learning, e.g., for evaluating the answers of short answer questions (cf. Pérez Marin, 2007; Ziai, 2018, and references therein) or in essay scoring (Shermis & Burstein, 2013). In terms of NLP methods, it also directly connects to the growing body of research on recognizing textual entailment and paraphrase recognition (Androutsopoulos & Malakasiotis, 2010; Dzikovska et al., 2013).

Shifting the focus from the analysis techniques to the interpretation of the analysis, just like the activity type and the learner play an important role in defining the space of variation to be dealt with by the NLP, the interpretation and feedback provided to the learner needs to be informed by activity and learner modelling (Amaral & Meurers, 2008). While feedback in human-computer interaction cannot simply be equated with that in human-human interaction, the results presented by Petersen (2010) for a dialogue-based ILTS indicate that results from the significant body of research on the effectiveness of different types of feedback in in-

structed SLA can transfer across modes, providing fresh momentum for research on feedback in the CALL domain (e.g., Pujolà, 2001).

A final important issue arising from the use of NLP in ILTS concerns the resulting lack of teacher autonomy. Quixal et al. (2010) explore putting the teacher back in charge of designing their activities with the help of an *ICALL authoring system* – a complex undertaking since in contrast to regular CALL authoring software, NLP analysis of learner language needs to be integrated without presupposing any understanding of the capabilities and limits of the NLP.

Learner Corpora

While there seems to have been little interaction between ILTS and learner corpus research, perhaps because ILTS traditionally have focused on exercises whereas most learner corpora consist of essays, the analysis of learner language in the annotation of learner corpora (cf. Granger, this volume) can be seen as an offline version of the online analysis performed by an ILTS (Meurers, 2015).

What motivates the use of NLP for learner corpora? In contrast to the automatic analysis of learner language in an ILTS providing feedback to the learner, the annotation of learner corpora essentially provides an index to learner language properties in support of the goal of advancing our understanding of acquisition in SLA research and to develop instructional methods and materials in FLT. Corpus annotation can support a more direct mapping from theoretical research questions to corpus instances (Meurers, 2005), yet for a reliable mapping it is crucial for corpus annotation to provide only reliable distinctions which are replicably based on the evidence found in the corpus and its meta-information, for which clear measures are available (Artstein & Poesio, 2009).

What is the nature of the learner language properties to be identified? Just as for ILTS, much of the work on annotating learner corpora has traditionally focused on learner errors, for which a number of *error annotation schemes* have been developed (cf. Díaz Negrillo & Fernández Domínguez, 2006, and references therein). There so far is no consensus, though, on the external and internal criteria, i.e., which error distinctions are needed for which purpose and which distinctions can reliably be annotated based on the evidence in the corpus and any meta-information available about the learner and the activity which the language was produced for. An explicit and reliable error annotation scheme and a gold standard reference corpus exemplifying it is an important next step for the development of automatic error annotation approaches, which need an agreed upon gold standard for development as well as for testing and comparison of approaches. Automating the currently manual error annotation process using NLP would support the annotation of significantly larger learner corpora and thus increase their usefulness for SLA research and FLT.

Since error annotation results from the annotator's comparison of a learner response to hypotheses about what the learner was trying to say, Lüdeling et al. (2005) argue for making the *target hypotheses* explicit in the annotation. This also makes it possible to specify alternative error annotations for the same sentence based on different target hypotheses in multi-level corpus annotation. However, Fitzpatrick & Seegmiller (2004) report unsatisfactory levels of inter-annotator agreement in determining such target hypotheses. It is an open research issue to determine for which type of learner responses written by which type of learners for which type of tasks such target hypotheses can reliably be determined. Target hypotheses might have to be limited to encoding only the minimal commitments necessary for error identification; and in place of target hypothesis strings, they might have to be formulated at more abstract levels, e.g., lemmas in topological fields, or sets of concepts the learner was trying to

express. In either case, if the target hypotheses are made explicit, the second step from target hypothesis to error identification can be studied separately and can be realized more reliably (Rosen et al., 2014).

Returning to the general question about the nature of the learner language properties which are relevant, SLA research essentially observes correlations of linguistic properties, whether erroneous or not. And even research focusing on learner errors needs to identify correlations with linguistic properties, e.g., to identify overuse/underuse of certain patterns, or measures of language development. While the use of NLP tools trained on the native language corpora is a useful starting point for providing a range of linguistic annotations, an important next step is to explore the creation of annotation schemes and methods capturing the *linguistic properties of learner language* (cf. Meurers & Dickinson, 2017, and references therein).

In terms of using NLP for providing *general measures of language development*, as captured by the triad Complexity, Accuracy, and Fluency (CAF, Housen & Kuiken, 2009), the automatic analysis of linguistic complexity in learner language has been a particularly active area of research (cf. Lu, 2014; Kyle, 2016; Chen, 2018, and references therein). The elaborateness and variedness of language use can be identified at all levels of linguistic modeling, including morphology, lexicon, syntax, and discourse.

To identify specific patterns that are characteristic of language development, it often is necessary to linguistically annotate the data (Meurers & Dickinson, 2017), which for large corpus resources requires automatic analysis. For example, Hawkins & Buttery (2009) identify so-called criterial features distinguishing different proficiency levels on the basis of part-of-speech tagged and parsed portions of the Cambridge Learner Corpus, and Alexopoulou et al. (2015) illustrate this with a study of relative clause development in the very large EFCamDat learner corpus (<https://corpus.mml.cam.ac.uk/efcamdat2>). Using the second release of that corpus, containing 1.2 million texts written by 175 thousand learners, Alexopoulou et al. (2017) highlight that the valid interpretation of linguistic complexity analysis also requires taking the properties of the task into account for which the writing was produced.

Learner corpora are also systematically analyzed with NLP methods to identify cross-linguistic effects, such as the transfer of characteristics of one's native language to text written in a second language. NLP research in this area was popularized by a series of shared tasks on Native Language Identification (Malmasi et al., 2017), with approaches exploring both shallow, surface-based and deep linguistic features (Jarvis & Crossley, 2012; Meurers et al., 2014; Bich, 2017). While most work focused on non-native English writing, Malmasi & Dras (2017) provide a genuinely multilingual approach.

2 NLP and the Analysis of Native Language for Learners

The second domain in which NLP connects to language learning derives from the need to expose learners to authentic, native language and its properties and to give them opportunity to interact with it. This includes work on searching for and enhancing authentic texts to be read by learners as well as the automatic generation of activities and tests from such texts. In contrast to the ILTS side of ICALL research covered in the first part of this article, the NLP in the applications under discussion here is used to process native language in Authentic Texts, hence referred to as ATICALL.

Most NLP research is developed and optimized for native language material and it is easier to obtain enough annotated language material to train the statistical models and machine learning approaches used in current research and development, so that in principle a wide

range of NLP tools with high quality analysis is available – even though this does not preempt the question which language properties are relevant for ATICALL applications and whether those properties can be derived from the ones targeted by standard NLP.

What motivates the use of NLP in ATICALL? Compared to using prefabricated materials such as those presented in textbooks, NLP-enhanced searching for materials in resources such as large corpora or the web makes it possible to provide on-demand, individualized access to up-to-date materials. It supports selecting and enhancing the presentation of texts depending on the background of a given learner, the specific contents of interest to them, and the language properties and forms of particular relevance given the sequencing of language materials appropriate for the learner’s stage (cf., e.g., Pienemann, 1998).

What is the nature of the learner language properties to be identified? Traditionally the most prominent property used for selecting texts has been a general notion of readability, for which a number of *readability formulas* were developed (DuBay, 2004). The traditional measures are based on shallow, easy to count features, typically average sentence and word lengths, but current machine learning methods informed by a broader range of linguistic characteristics are substantially more accurate (cf., e.g., Xia et al., 2016; Crossley et al., 2017; Weiss & Meurers, 2018), including some commercial systems (Nelson et al., 2012). Interestingly, readability analysis substantially benefits from the integration of complexity features originally designed to measure second language development (Vajjala & Meurers, 2012). Work in the Coh-Metrix project emphasizes the importance of analyzing text cohesion and coherence and of taking a reader’s cognitive aptitudes into account for making predictions about reading comprehension (McNamara et al., 2014). For languages other than English, morphological features become more prominent (Dell’Orletta et al., 2011; François & Fairon, 2012; Hancke et al., 2012), for which finite-state NLP approaches provide particularly rich information (Reynolds, 2016), also readily supporting exercise generation on that basis (Antonsen & Argese, 2018). For reading practice with a focus on vocabulary acquisition (Cobb, 2008), several projects have emphasized the relevance and impact of individual learner models (Walmsley, 2015; Heilman et al., 2010).

Going beyond readability, based on the insight from SLA research that awareness of language categories and forms is an important ingredient for successful second language acquisition (cf. Lightbown & Spada, 1999), a wide range of linguistic properties have been identified as relevant for language awareness, including morphological, syntactic, semantic, and pragmatic information (cf. Schmidt, 1995, p. 30). In response to this need, the FLAIR system (Chinkina & Meurers, 2016) supports linguistically-aware web search, which makes it possible to systematically enrich the input of language learners with the kind of language patterns to be acquired next. By integrating automated linguistic complexity analysis, it also becomes possible to retrieve reading material in the zone of proximal development of a learner by matching the complexity of the material to that of text written by this learner (Chen & Meurers, 2019).

Complementing the question of how to obtain material for language learners, there are several strands of ATICALL applications which focus on the enhanced presentation of and learner interaction with such materials. One group of NLP-based tools such as COMPASS (Breidt & Feldweg, 1997), Glosser (Nerbonne et al., 1998) and Grim (Knutsson et al., 2007) provides a ***reading environment*** in which texts in a foreign language can be read with quick access to dictionaries, morphological information, and concordances. The Alpheios project (<http://alpheios.net>) focuses on literature in Latin and ancient Greek, providing links between words and translations, access to online grammar reference, and a quiz mode asking the learner to identify which word corresponds to which translation.

Another strand of ATICALL research focuses on supporting language awareness by *automating input enhancement* (Sharwood Smith, 1993), i.e., by realizing strategies which highlight the salience of particular language categories and forms. For instance, WERTi (Meurers et al., 2010) visually enhances web pages and automatically generates activities for language patterns which are known to be difficult for learners of English, such as determiners and prepositions, phrasal verbs, the distinction between gerunds and to-infinitives, and *wh*-question formation. Complementing the visual input enhancement of forms, Chinkina & Meurers (2017) propose to automatically generate questions as a form of functionally-driven input enhancement. One can view such automatic input enhancement as an enrichment of Data-Driven Learning (DDL). Where DDL has been characterized as an “attempt to cut out the middleman [the teacher] as far as possible and to give the learner direct access to the data” (Boulton 2009, p. 82, citing Tim Johns), in visual input enhancement the learner stays in control, but the NLP uses ‘teacher knowledge’ about relevant and difficult language properties to make those more prominent and noticeable for the learner.

A final, prominent strand of NLP research in this domain addresses the *generation of exercises and tests*. Most of the work has targeted the automatic generation of multiple choice cloze tests for language assessment and vocabulary drill (cf., e.g. Sumita et al., 2005; Liu et al., 2005, and references therein). Issues involving NLP in this domain include the selection of seed sentences, the determination of appropriate blank positions, and the generation of good distractor items. The VISL project (Bick, 2005) also includes a tool supporting the generation of automatic cloze exercises, which is part of an extensive NLP-based environment of games and corpus tools aimed at fostering linguistic awareness for dozens of languages (<http://beta.visl.sdu.dk>). Finally, the Task Generator for ESL (Toole & Heift, 2001) supports the creation of gap-filling and build-a-sentence exercises such as the ones found in an ILTS. The instructor provides a text, chooses from a list of learning objectives (e.g., plural, passive), and select the exercise type to be generated. The Task Generator supports complex language patterns and provides formative feedback based on NLP analysis of learner responses, bringing us full circle to the research on ILTS we started with. Currently the most advanced approach in this line of research is the Language Muse Activity Palette (Burstein & Sabatini, 2016).

In conclusion, the use of NLP in the context of learning language offers rich opportunities, both in terms of developing applications in support of language teaching and learning and in terms of supporting SLA research – even though NLP so far has only had limited impact on real-life language teaching and SLA. More interdisciplinary collaboration between SLA and NLP will be crucial for developing reliable annotation schemes and analysis techniques which identify the properties which are relevant and important for analyzing learner language and analyzing language for learners.

SEE ALSO: Automatic Speech Recognition; Computer Assisted Pronunciation Teaching; Computer-Assisted Language Learning Effectiveness Research; Corpus Linguistics in Language Teaching; Innovation in Language Teaching and Learning; Learner Corpora; Mobile Assisted Language Learning;

References

Alexopoulou, T., J. Geertzen, A. Korhonen & D. Meurers (2015). Exploring Big Educational Learner Corpora for SLA Research: Perspectives on Relative Clauses. *International Journal of Learner Corpus Research* 1(1), 96–129.

- Alexopoulou, T., M. Michel, A. Murakami & D. Meurers (2017). Task Effects on Linguistic Complexity and Accuracy: A Large-Scale Learner Corpus Analysis Employing Natural Language Processing Techniques. *Language Learning* 67, 181–209.
- Amaral, L. & D. Meurers (2008). From Recording Linguistic Competence to Supporting Inferences about Language Acquisition in Context: Extending the Conceptualization of Student Models for Intelligent Computer-Assisted Language Learning. *Computer-Assisted Language Learning* 21(4), 323–338.
- Amaral, L. & D. Meurers (2011). On Using Intelligent Computer-Assisted Language Learning in Real-Life Foreign Language Teaching and Learning. *ReCALL* 23(1), 4–24.
- Androutsopoulos, I. & P. Malakasiotis (2010). A Survey of Paraphrasing and Textual Entailment Methods. *Journal of Artificial Intelligence Research* 38, 135–187.
- Antonsen, L. & C. Argese (2018). Using authentic texts for grammar exercises for a minority language. In *Proceedings of the 7th workshop on NLP for Computer Assisted Language Learning*. Stockholm, Sweden, pp. 1–9. URL <http://aclweb.org/anthology/W18-7101>.
- Artstein, R. & M. Poesio (2009). Survey Article: Inter-Coder Agreement for Computational Linguistics. *Computational Linguistics* 34(4), 1–42.
- Bachman, L. F. & A. S. Palmer (1996). *Language Testing in Practice: Designing and Developing Useful Language Tests*. Oxford University Press.
- Bailey, S. & D. Meurers (2008). Diagnosing meaning errors in short answers to reading comprehension questions. In *Proceedings of the 3rd Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*. Columbus, Ohio, pp. 107–115. URL <http://aclweb.org/anthology/W08-0913>.
- Baker, D. (1989). *Language Testing: A Critical Survey and Practical Guide*. London: Edward Arnold.
- Bich, S. (2017). "Is There Choice in Non-Native Voice?" Linguistic Feature Engineering and a Variationist Perspective in Automatic Native Language Identification. Ph.D. thesis, Eberhard-Karls Universität Tübingen. URL <http://hdl.handle.net/10900/77443>.
- Bick, E. (2005). Grammar for Fun: IT-based Grammar Learning with VISL. In P. Juel (ed.), *CALL for the Nordic Languages*, Copenhagen: Samfundslitteratur, Copenhagen Studies in Language, pp. 49–64. URL <http://beta.visl.sdu.dk/pdf/CALL2004.pdf>.
- Boulton, A. (2009). Data-driven Learning: Reasonable Fears and Rational Reassurance. *Indian Journal of Applied Linguistics* 35(1), 81–106.
- Boyd, A. (2012). Detecting and Diagnosing Grammatical Errors for Beginning Learners of German: From Learner Corpus Annotation to Constraint Satisfaction Problems. Ph.D. thesis, The Ohio State University. URL http://rave.ohiolink.edu/etdc/view?acc_num=osu1325170396.
- Breidt, E. & H. Feldweg (1997). Accessing Foreign Languages with COMPASS. *Machine Translation* 12(1–2), 153–174.
- Burstein, J. & J. Sabatini (2016). The Language Muse Activity Palette. In S. A. Crossley & D. S. McNamara (eds.), *Adaptive Educational Technologies for Literacy Instruction*, Routledge, pp. 275–280.
- Chen, X. (2018). Automatic Analysis of Linguistic Complexity and Its Application in Language Learning Research. Ph.D. thesis, Eberhard Karls Universität Tübingen Germany. URL <http://hdl.handle.net/10900/85888>.

- Chen, X. & D. Meurers (2019). Linking text readability and learner proficiency using linguistic complexity feature vector distance. *Computer-Assisted Language Learning* URL <https://doi.org/10.1080/09588221.2018.1527358>.
- Chinkina, M. & D. Meurers (2016). Linguistically-Aware Information Retrieval: Providing Input Enrichment for Second Language Learners. In *Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*. San Diego, CA: ACL, pp. 188–198. URL <http://aclweb.org/anthology/W16-0521.pdf>.
- Chinkina, M. & D. Meurers (2017). Question Generation for Language Learning: From ensuring texts are read to supporting learning. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*. Copenhagen, Denmark, pp. 334–344. URL <http://aclweb.org/anthology/W17-5038.pdf>.
- Choi, I.-C. (2016). Efficacy of an ICALL tutoring system and process-oriented corrective feedback. *Computer Assisted Language Learning* 29(2), 334–364.
- Chollampatt, S. & H. T. Ng (2018). A multilayer convolutional encoder-decoder neural network for grammatical error correction. In *Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*. New Orleans, pp. 5755–5762.
- Cobb, T. (2008). Necessary or Nice? The Role of Computers in L2 Reading. In Z. Han & N. J. Anderson (eds.), *L2 Reading Research and Instruction*, Ann Arbor: University of Michigan Press, pp. 144–172.
- Crossley, S. A., S. Skalicky, M. Dascalu, D. S. McNamara & K. Kyle (2017). Predicting text comprehension, processing, and familiarity in adult readers: New approaches to readability formulas. *Discourse Processes* 54(5-6), 340–359.
- De Felice, R. (2008). Automatic Error Detection in Non-native English. Ph.D. thesis, St Catherine's College, University of Oxford. URL <http://purl.org/net/DeFelice-08.pdf>.
- Dell'Orletta, F., S. Montemagni & G. Venturi (2011). READ-IT: Assessing Readability of Italian Texts with a View to Text Simplification. In *Proceedings of the 2nd Workshop on Speech and Language Processing for Assistive Technologies*. pp. 73–83.
- Delmonte, R. (2003). Linguistic knowledge and reasoning for error diagnosis and feedback generation. *CALICO Journal* 20(3), 513–532.
- Díaz Negrillo, A. & J. Fernández Domínguez (2006). Error Tagging Systems for Learner Corpora. *Revista Española de Lingüística Aplicada (RESLA)* 19, 83–102.
- Dickinson, M. (2006). Writer's Aids. In K. Brown (ed.), *Encyclopedia of Language and Linguistics*, Oxford: Elsevier. 2 ed.
- DuBay, W. H. (2004). *The Principles of Readability*. Costa Mesa, California: Impact Information. URL <http://impact-information.com/impactinfo/readability02.pdf>.
- Dzikovska, M., R. Nielsen et al. (2013). SemEval-2013 Task 7: The Joint Student Response Analysis and 8th Recognizing Textual Entailment Challenge. In *Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*. Atlanta, Georgia, USA, pp. 263–274. URL <http://aclweb.org/anthology/S13-2045>.
- Ellis, R. (2009). Task-based language teaching: sorting out the misunderstandings. *International Journal of Applied Linguistics* 19(3), 221–246.
- Fitzpatrick, E. & M. S. Seegmiller (2004). The Montclair electronic language database project. In U. Connor & T. Upton (eds.), *Applied Corpus Linguistics: A Multidimensional Perspective*, Amsterdam: Rodopi.

- Foth, K., W. Menzel & I. Schröder (2005). Robust parsing with weighted constraints. *Natural Language Engineering* 11(1), 1–25.
- François, T. & C. Fairon (2012). 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*. pp. 466–477.
- Gamon, M., C. Leacock, C. Brockett, W. B. Dolan, J. Gao, D. Belenko & A. Klementiev (2009). Using Statistical Techniques and Web Search to Correct ESL Errors. *CALICO Journal* 26(3), 491–511. URL <http://purl.org/calico/Gamon.Leacock.ea-09.pdf>.
- Granger, S. (this volume). Learner Corpora. In C. A. Chapelle (ed.), *The Concise Encyclopedia of Applied Linguistics*, Wiley.
- Hancke, J., S. Vajjala & D. Meurers (2012). Readability Classification for German using lexical, syntactic, and morphological features. In *Proceedings of the 24th International Conference on Computational Linguistics (COLING)*. Mumbai, India, pp. 1063–1080. URL <http://aclweb.org/anthology/C12-1065>.
- Hawkins, J. A. & P. Buttery (2009). Using Learner Language from Corpora to Profile Levels of Proficiency – Insights from the English Profile Programme. In *Studies in Language Testing: The Social and Educational Impact of Language Assessment*, Cambridge: Cambridge University Press.
- Heift, T. (2010). Developing an Intelligent Language Tutor. *CALICO Journal* 27(3), 443–459.
- Heift, T. (2017). History and key developments in intelligent computer-assisted language learning (ICALL). In S. L. Thorne & S. May (eds.), *Language, education and technology*, Springer, pp. 289–312. Third edition ed.
- Heift, T. & M. Schulze (2007). *Errors and Intelligence in Computer-Assisted Language Learning: Parsers and Pedagogues*. New York: Routledge.
- Heift, T. & M. Schulze (2015). Tutorial computer-assisted language learning. *Language Teaching* 48(4), 471–490.
- Heilman, M., K. Collins-Thompson, J. Callan, M. Eskenazi, A. Juffs & L. Wilson (2010). Personalization of Reading Passages Improves Vocabulary Acquisition. *International Journal of Artificial Intelligence in Education* 20, 73–98.
- Housen, A. & F. Kuiken (2009). Complexity, Accuracy and Fluency in Second Language Acquisition. *Applied Linguistics* 30(4), 461–473.
- Jarvis, S. & S. A. Crossley (eds.) (2012). *Approaching Language Transfer through Text Classification: Explorations in the Detection-based Approach*. Second Language Acquisition. Multilingual Matters. URL http://books.google.de/books?id=dyH0_DPdQV0C.
- Johns, T. (1994). From printout to handout: Grammar and vocabulary teaching in the context of data-driven learning. In T. Odlin (ed.), *Perspectives on Pedagogical Grammar*, Cambridge: Cambridge University Press, pp. 293–313.
- Johnson, M. (1994). Two ways of formalizing grammars. *Linguistics and Philosophy* 17(3), 221–248.
- Junczys-Dowmunt, M. & R. Grundkiewicz (2016). Phrase-based Machine Translation is State-of-the-Art for Automatic Grammatical Error Correction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, pp. 1546–1556. URL <http://aclweb.org/anthology/>

D16-1161.

- Jurafsky, D. & J. H. Martin (2009). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Upper Saddle River, NJ: Prentice Hall, 2nd edition ed.
- Knutsson, O., T. C. Pargman, K. S. Eklundh & S. Westlund (2007). Designing and developing a language environment for second language writers. *Computers & Education* 49(4), 1122 – 1146.
- Kwasny, S. C. & N. K. Sondheimer (1981). Relaxation Techniques for Parsing Grammatically Ill-Formed Input in Natural Language Understanding Systems. *American Journal of Computational Linguistics* 7(2), 99–108. URL <http://purl.org/net/Kwasny.Sondheimer-81.pdf>.
- Kyle, K. (2016). Measuring Syntactic Development in L2 Writing: Fine Grained Indices of Syntactic Complexity and Usage-Based Indices of Syntactic Sophistication. Ph.D. thesis, Georgia State University. URL http://scholarworks.gsu.edu/alesl_diss/35.
- Lightbown, P. M. & N. Spada (1999). *How languages are learned*. Oxford: Oxford University Press. URL http://books.google.com/books?id=w1YTbuCsR7wC&lpg=PR11&ots=_DdFbrMCOO&dq=How%20languages%20are%20learned&lr&pg=PR11#v=onepage&q&f=false.
- Liu, C.-L., C.-H. Wang, Z.-M. Gao & S.-M. Huang (2005). Applications of Lexical Information for Algorithmically Composing Multiple-Choice Cloze Items. In *Proceedings of the Second Workshop on Building Educational Applications Using NLP*. pp. 1–8. URL <http://aclweb.org/anthology/W05-0201>.
- Lu, X. (2014). *Computational Methods for Corpus Annotation and Analysis*. Springer.
- Lüdeling, A., M. Walter, E. Kroymann & P. Adolphs (2005). Multi-level error annotation in learner corpora. In *Proceedings of Corpus Linguistics*. Birmingham. URL <http://www.corpus.bham.ac.uk/PCLC/Falko-CL2006.doc>.
- Malmasi, S. & M. Dras (2017). Multilingual native language identification. *Natural Language Engineering* 23(2), 163–215.
- Malmasi, S., K. Evanini, A. Cahill, J. Tetreault, R. Pugh, C. Hamill, D. Napolitano & Y. Qian (2017). A Report on the 2017 Native Language Identification Shared Task. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*. pp. 62–75.
- Matthews, C. (1992). Going AI: Foundations of ICALL. *Computer Assisted Language Learning* 5(1), 13–31.
- McNamara, D. S., A. C. Graesser, P. M. McCarthy & Z. Cai (2014). *Automated evaluation of text and discourse with Coh-Metrix*. Cambridge University Press.
- Meurers, D. (2005). On the use of electronic corpora for theoretical linguistics. Case studies from the syntax of German. *Lingua* 115(11), 1619–1639.
- Meurers, D. (2015). Learner Corpora and Natural Language Processing. In S. Granger, G. Gilquin & F. Meunier (eds.), *The Cambridge Handbook of Learner Corpus Research*, Cambridge University Press, pp. 537–566. <http://purl.org/dm/papers/meurers-15.html>.
- Meurers, D. & M. Dickinson (2017). Evidence and Interpretation in Language Learning Research: Opportunities for Collaboration with Computational Linguistics. *Language Learning* 67(2).

- Meurers, D., J. Krivanek & S. Bykh (2014). On the Automatic Analysis of Learner Corpora: Native Language Identification as Experimental Testbed of Language Modeling between Surface Features and Linguistic Abstraction. In A. Alcaraz Sintes & S. Valera (eds.), *Diachrony and Synchrony in English Corpus Studies*. Frankfurt am Main: Peter Lang, pp. 285–314.
- Meurers, D., K. D. Kuthy, F. Nuxoll, B. Rudzewitz & R. Ziai (2019). Scaling up intervention studies to investigate real-life foreign language learning in school. *Annual Review of Applied Linguistics* 39. To appear.
- Meurers, D., R. Ziai, L. Amaral, A. Boyd, A. Dimitrov, V. Metcalf & N. Ott (2010). Enhancing Authentic Web Pages for Language Learners. In *Proceedings of the 5th Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*. Los Angeles: ACL, pp. 10–18. URL <http://aclweb.org/anthology/W10-1002.pdf>.
- Michaud, L. N. & K. F. McCoy (2004). Empirical Derivation of a Sequence of User Stereotypes for Language Learning. *User Modeling and User-Adapted Interaction* 14(4), 317–350.
- Naber, D. (2003). A Rule-Based Style and Grammar Checker. Master's thesis, Universität Bielefeld. URL <http://www.danielnaber.de/publications>.
- Nagata, N. (2009). Robo-Sensei's NLP-Based Error Detection and Feedback Generation. *CALICO Journal* 26(3), 562–579. URL <http://purl.org/calico/nagata09.pdf>.
- Nelson, J., C. Perfetti, D. Liben & M. Liben (2012). *Measures of Text Difficulty: Testing their Predictive Value for Grade Levels and Student Performance*. Tech. rep., The Council of Chief State School Officers. URL <http://purl.org/net/Nelson.Perfetti.ea-12.pdf>.
- Nerbonne, J. (2003). Natural language processing in computer-assisted language learning. In R. Mitkov (ed.), *The Oxford Handbook of Computational Linguistics*, Oxford University Press.
- Nerbonne, J., D. Dokter & P. Smit (1998). Morphological Processing and Computer-Assisted Language Learning. *Computer Assisted Language Learning* 11(5), 543–559.
- Pennington, M. C. & P. Rogerson-Revell (2019). Using Technology for Pronunciation Teaching, Learning, and Assessment. In *English Pronunciation Teaching and Research*, Springer, pp. 235–286.
- Pérez Marin, D. R. (2007). Adaptive Computer Assisted Assessment of free-text students' answers: an approach to automatically generate students' conceptual models. Ph.D. thesis, Universidad Autonoma de Madrid.
- Petersen, K. (2010). Implicit Corrective Feedback in Computer-Guided Interaction: Does Mode Matter? Ph.D. thesis, Georgetown University. URL <http://hdl.handle.net/10822/553155>.
- Pienemann, M. (1998). *Language Processing and Second Language Development: Processability Theory*. Amsterdam: John Benjamins.
- Pujolà, J.-T. (2001). Did CALL Feedback Feed Back? Researching Learners' Use of Feedback. *ReCALL* 13(1), 79–98.
- Quixal, M. & D. Meurers (2016). How can writing tasks be characterized in a way serving pedagogical goals and automatic analysis needs? *CALICO Journal* 33, 19–48. URL <http://purl.org/dm/papers/Quixal.Meurers-16.html>.
- Quixal, M., S. Preuß, B. Boullosa & D. García-Narbona (2010). AutoLearn's authoring tool:

- a piece of cake for teachers. In *Proceedings of the NAACL HLT 2010 Fifth Workshop on Innovative Use of NLP for Building Educational Applications*. Los Angeles, pp. 19–27. URL <http://aclweb.org/anthology/W10-1003.pdf>.
- Ramsay, A. & V. Mirzaiean (2005). Content-based support for Persian learners of English. *ReCALL* 17, 139–154.
- Reuer, V. (2003). Error recognition and feedback with Lexical Functional Grammar. *CALICO Journal* 20(3), 497–512. URL <http://purl.org/calico/Reuer-03.pdf>.
- Reynolds, R. (2016). Russian natural language processing for computer-assisted language learning: capturing the benefits of deep morphological analysis in real-life applications. Ph.D. thesis, UiT - The Arctic University of Norway. <https://munin.uit.no/handle/10037/9685>.
- Rimrott, A. & T. Heift (2008). Evaluating automatic detection of misspellings in German. *Language Learning and Technology* 12(3), 73–92. URL <http://llt.msu.edu/vol12num3/rimrottheift.pdf>.
- Rosen, A., J. Hana, B. Štindlová & A. Feldman (2014). Evaluating and automating the annotation of a learner corpus. *Language Resources and Evaluation* 48(1), 65–92.
- Schmidt, R. (1995). Consciousness and foreign language learning: A tutorial on the role of attention and awareness in learning. In R. Schmidt (ed.), *Attention and awareness in foreign language learning*, Honolulu, HI: University of Hawaii, pp. 1–63.
- Schulze, M. (2011). Learner Modeling in Intelligent Computer-Assisted Language Learning. In C. A. Chapelle (ed.), *The Encyclopedia of Applied Linguistics. Technology and Language*, Wiley-Blackwell.
- Sharwood Smith, M. (1993). Input enhancement in instructed SLA: Theoretical bases. *Studies in Second Language Acquisition* 15, 165–179. URL <https://doi.org/10.1017/s0272263100011943>.
- Shermis, M. D. & J. Burstein (eds.) (2013). *Handbook on Automated Essay Evaluation: Current Applications and New Directions*. London and New York: Routledge, Taylor & Francis Group.
- Sleeman, D. (1982). Inferring (mal) rules from pupil's protocols. In *Proceedings of the 5th European Conference on Artificial Intelligence (ECAI)*. Orsay, France, pp. 160–164.
- Sumita, E., F. Sugaya & S. Yamamoto (2005). Measuring Non-native Speakers' Proficiency of English by Using a Test with Automatically-Generated Fill-in-the-Blank Questions. In *Proceedings of the Second Workshop on Building Educational Applications Using NLP*. pp. 61–68. URL <http://aclweb.org/anthology/W05-0210>.
- Toole, J. & T. Heift (2001). Generating Learning Content for an Intelligent Language Tutoring System. In *Proceedings of NLP-CALL Workshop at the 10th Int. Conf. on Artificial Intelligence in Education (AI-ED)*. San Antonio, Texas, pp. 1–8.
- Tschichold, C. (1999). Grammar checking for CALL: Strategies for improving foreign language grammar checkers. In K. Cameron (ed.), *CALL: Media, Design & Applications*, Exton, PA: Swets & Zeitlinger.
- Vajjala, S. & D. Meurers (2012). On Improving the Accuracy of Readability Classification using Insights from Second Language Acquisition. In *Proceedings of the 7th Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*. Montréal, Canada: ACL, pp. 163–173. URL <http://aclweb.org/anthology/W12-2019.pdf>.

- Wagner, J. & J. Foster (2009). The effect of correcting grammatical errors on parse probabilities. In *Proceedings of the 11th International Conference on Parsing Technologies (IWPT'09)*. Paris, France: Association for Computational Linguistics, pp. 176–179. URL <http://aclweb.org/anthology/W09-3827>.
- Walmsley, M. (2015). Learner Modelling for Individualised Reading in a Second Language. Ph.D. thesis, The University of Waikato. URL <http://hdl.handle.net/10289/10559>.
- Weischedel, R. M. & N. K. Sondheimer (1983). Meta-rules as a Basis for Processing Ill-formed Input. *Computational Linguistics* 9(3-4), 161–177. URL <http://aclweb.org/anthology/J83-3003>.
- Weiss, Z. & D. Meurers (2018). Modeling the Readability of German Targeting Adults and Children: An Empirically Broad Analysis and its Cross-Corpus Validation. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*. Santa Fe, New Mexico, USA: International Committee on Computational Linguistic.
- Xia, M., E. Kochmar & T. Briscoe (2016). Text readability assessment for second language learners. In *Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications*. pp. 12–22.
- Ziai, R. (2018). Short Answer Assessment in Context: The Role of Information Structure. Ph.D. thesis, Eberhard-Karls Universität Tübingen. URL <http://hdl.handle.net/10900/81732>.

Suggested Readings

Rebuschat, P., Detmar, M., & McEnery, T. (eds.) (2017). Language learning research at the intersection of experimental, computational and corpus-based approaches. *Language Learning*, 67(S1). <https://doi.org/10.1111/lang.12243>

The Proceedings of the Workshop Series on *Innovative Use of NLP for Building Educational Applications (BEA)*, organized by the Special Interest Group (SIG) for Building Educational Applications of the Association for Computational Linguistics, accessible from <https://sig-edu.org/bea/past>

The Proceedings of the Workshop Series *Natural Language Processing for Computer-Assisted Language Learning (NLP4CALL)* organized by the ICALL SIG of the North European Association of Language Technology, accessible from <https://spraakbanken.gu.se/eng/research/icall/nlp4call>