

Enhancing a Web-based Language Tutoring System with Learning Analytics

Björn Rudzewitz, Ramon Ziai, Florian Nuxoll, Kordula De Kuthy, Detmar Meurers
ICALL Research Group & SFB 833 & LEAD
University of Tübingen, Germany
{br,rzai,fnuxoll,kdk,dm}@sfs.uni-tuebingen.de

ABSTRACT

This paper develops learning analytics and educational data mining perspectives for the FeedBook, an Intelligent Language Tutoring System that was fully integrated as a homework platform into 14 regular 7th grade English classes in German secondary schools during a full-year study ([12]). We demonstrate how different perspectives on learner and interaction data supported by different user interfaces of the system can help address demands of students, teachers, and material designers. The interfaces offers perspectives on individual students in the form of an open learner model, on classes and their typical errors to inform teachers, and on general learner performance to provide insights into the effectiveness of exercises for material designers.

1. INTRODUCTION AND CONTEXT

The digital transformation of many areas of society, including education, changes traditional roles, tasks, and workflows. Romero and Ventura ([17]) point out that digitalization in the education context correlates with the rise of web-based information and technologies. Web-based learning fosters processes where teachers and students are, on the one hand, more distant from each other, in that they can work from anywhere and are not tied to a physical classroom. This leads to a loss of information since teachers cannot monitor students directly. On the other hand, this development also leads to an increase in educational data available since every interaction with a web-based system generates data points that can be stored and analyzed. Correspondingly, in web-based learning, the role and tasks of teachers are changed. Such systems can take over certain tasks traditionally performed by humans, such as monitoring students, correcting exercises, or aggregating information on errors.

For students, the key difference when working digitally compared to on paper is the nature of the interaction with the material. They can work from anywhere at any time on any material and, depending on system functionality, can receive immediate feedback on their output. Such interaction can

lead to students revising misconceptions and producing answers that often contain fewer errors compared to homework done on paper. At the same time, teachers need information about common misconceptions of their students in order to be able to prepare their classroom teaching. They need some insights into the process of the interaction of their students, the steps performed while working on an exercise, individually or in aggregated form. The collection of this type of data creates large quantities of learner data that would be unavailable without digitalization.

To deal with the increased amount of data in a digital learning context potentially allowing an increased distance between teachers and students, it is necessary to develop and provide pedagogically informed methods and tools that use the available data and provide interfaces addressing real-life needs for all stakeholders in the educational context. Learning Analytics (LA) as “the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimizing learning and the environments in which it occurs” ([21]) and Educational Data Mining (EDM) “developing methods for exploring the unique types of data that come from educational settings, using those methods to better understand students, and the settings which they learn in” ([6]) can help address this need for pedagogically informed methods for the analysis of educational data. LA and EDM methods presuppose the availability of learning process and product data that can be analyzed. While such data can in principle be collected in any context, most commonly they originate from web-based learning platforms given that in such a setup, all data can readily be collected in real-time, stored centrally on a server, and be presented in different ways to different users of the system at any time.

Extensive research has been conducted on developing tutoring systems in the mathematical domain (cf., e.g., [20]). In such formal domains, there is a one-to-one correspondence between the form and the meaning of student productions. On the other hand, language learning (and the use of language in support of content learning) is sometimes characterized as an “ill-defined domain” ([9, 22]), where it is difficult to systematically characterize when an answer is an acceptable solution given that language supports a substantial number of well-formed paraphrases and ill-formed variability, requiring dedicated natural language processing techniques and an understanding of language learning to support valid interpretations of learner utterances ([11, 13]).

Second Language Acquisition (SLA) research attempts to characterize the nature of the language learning process. Interactionist approaches in SLA emphasize that for language learning, not only input is necessary, but also interaction ([8]) given a specific social configuration ([23]). Such interaction in an educational context typically involves peer learners and teachers. In Computer-Assisted Language Learning (CALL), digital tutoring systems may also support dyadic interaction, with some SLA studies showing comparable learning gains for human-computer as for human-human interaction (e.g., [16]). When learners do not successfully complete a task, the teacher or system can provide corrective feedback. The reaction to corrective feedback by learners is typically referred to as uptake ([10]). Uptake is considered successful if the learner manages to implement the specific feedback provided by the teacher, i.e., when the learner answer is correct or does not contain this particular error anymore. Explicit meta-linguistic corrective feedback has been shown to be an effective form of feedback supporting successful uptake (cf., e.g., [7]).

Bringing together research on SLA and CALL with real-life contexts of teaching and learning, we would like to argue that LA is essential to address the real-life needs of students, teachers, and material designers. In this paper, we substantiate this perspective by first introducing the intelligent web-based tutoring system FeedBook, which we used in a yearlong study in an authentic secondary school context, before turning to the LA interfaces and insights this supports in section 3.

2. THE BASIS FOR DATA COLLECTION

In order to collect data for LA and EDM, two components are necessary: an instrument to collect data with, and users who interact with this instrument, ideally in an authentic learning context to support ecologically valid, generalizable interpretations. In the following, we will first describe the instrument (a web-based tutoring system) and then the study we conducted in a real-life school context to generate the data underlying the work presented in this article.

2.1 The FeedBook System

The FeedBook ([18, 24, 14]) is a web-based tutoring system for 7th grade learners of English. The system offers a digital version of 230 exercises from the paper-based *Camden Town 3* workbook, and 154 exercises in three difficulty levels as additional material, discussed more in section 3.1. The FeedBook follows the paper workbook in its design, splitting the curriculum into six themes. Each theme has a specific topic, builds up to functional language tasks in the spirit of Task-Based Language Teaching ([5]), and covers two to three grammar constructs. For each of the grammar constructs in the curriculum, the system offers immediate interactive meta-linguistic formative scaffolding feedback to students ([19]). While students are working on an exercise, the system checks an answer immediately after it has been typed in and provides feedback on whether it is correct or *why* it is incorrect. As spelled out in [19], generation of the feedback combines computational linguistic analysis of exercise specifications with explicit modeling of the language constructs in the curriculum and potential misconceptions of language learners connected to these constructs. The FeedBook models 188 different error types and thereby covers

The screenshot shows a lesson titled 'Off to Greece again' with a task to compare two airlines, Midair and Air-Con. The feedback window highlights the error 'The tickets at Air-Con are expensiver than at Midair.' and provides the rule: 'When an adjective has three or more syllables, we form the comparative with "more" and the superlative with "most".' The feedback also indicates that 'The tickets at Air-Con are **expensiver** than at Midair.'

Figure 1: Exemplifying feedback in the FeedBook

all grammar topics in the curriculum of 7th grade English. Figure 1 shows an example for grammar feedback displayed to a learner, informing the learner about the type of error, a strategy to repair it, and the location of the error (which is displayed after clicking the magnifying glass in the bottom-left corner of the feedback window). We here focus on the perspective of the student – for teachers, the system also offers functions to inspect and correct the submissions of their students. We describe LA-related functions of the FeedBook in section 3.

2.2 Effectiveness Study

The FeedBook system has been used as a drop-in replacement for the printed workbook in fourteen 7th grade school classes since the beginning of this school year in September 2018 as part of a yearlong study ([12]). The goal of this study is to test the effectiveness of interactive scaffolding feedback targeting specific grammar constructs. As part of the study, the students in each class were randomly assigned to one of two groups. Both groups receive feedback on meaning, orthography, and default feedback (which asks the user to try again in case no known misconception was detected). The grammar constructions covered by the curriculum are divided into two, with students in one group receiving specific grammar feedback on one half of the grammar constructions in the curriculum, and the other on the other half. In practice, each group is associated with a blacklist of grammatical constructions. When a learner makes a grammar error that is associated with the learner's blacklist, the system does not display the specific grammar error. The only change to the classes was the introduction of the FeedBook in place of the printed workbook, and a request to the teachers to assign at least two to three homework exercises for each grammar topic they covered. For all students, including those where feedback is disabled for a specific grammar topic, teachers were also able to provide manual feedback in the system (like in a traditional paper-based setup).

Given that every theme is associated with two to three grammar constructs, which are assigned to the same group, it becomes possible to test the impact of the specific grammar

feedback before and after each theme. The first complete analysis we conducted ([12]) is based on the performance of 205 students working through the second theme of the workbook, using a pretest/posttest design. The grammar topics focused on in that theme are comparatives, conditional clauses, and relative clauses. The analysis shows that the learners who received the specific scaffolded feedback by the system for those grammar topics improved significantly more on the posttest, learning 62% more, compared to those who did not (Cohen's $d = 0.56$). The study thus supports the effectiveness of a computer-based interactive scaffolding system for language learning with substantial ecological validity given the fully authentic school setting.

3. LEARNING ANALYTICS IN FEEDBOOK

The primary goal of LA is to enable humans to make informed decisions when presented with educational data ([21]). Visualization plays an important role in making data interpretable. For this reason, we describe a subset of the LA functionality implemented in the FeedBook and currently in use by students, teachers, and material designers. We motivate the implementation of each interface by stating which questions the interface answers for the respective user group.

3.1 Student Perspective on Learner Model

For students, the most important questions in the educational context are: What have I learned? How well have I mastered a concept? Where should I go next to overcome misconceptions, address gaps in knowledge, or progress through new material?

In order to address these needs, we implemented an open learner model. The learner model is accessible to students from the start page of the system and provides information about the progress on each of the grammar topics to be acquired according to the 7th grade curriculum. Open learner models, in contrast to closed learner models, present to learners the learning-related information the system has collected about them. Reasons for opening learner models up to the learner are that this fosters the meta-cognition of learners. It gives learners insights into their learning process, supports them in diagnosing areas requiring additional practice, allows individualized navigation in the learning system via suggestions of material, and implements the right of users to know what a system knows about them ([2, 3]).

In order to not overwhelm the user with information, the learner model is organized in a hierarchical way. On the top level, it presents the grammar categories from the curriculum, such as *present tense* or *conditionals*. Upon clicking on such a category, the system displays a wind rose chart that, for every specific grammar construction in the selected category, plots the number of times the student used the construction correctly in green and the number of times it was not correctly used but demanded by the task in red. Figure 2 illustrates this. For every prompt, we know from the computational linguistic analysis of the task specification used in the feedback generation mechanism which grammar construct has to be mastered, in order to answer the prompt correctly. So we can count the number of times a student used a specific grammar construct correctly by looking at which prompts they answered correctly, either at first try, after feedback, or when submitting it to the teacher. For

Lernermodell betrachten

Hier werden Statistiken zum Lernverhalten angezeigt. Bitte unten auf eine Kategorie klicken um mehr zu erfahren.

future tense

present tense

Richtig vs. Fehler



simple present (regular verb)

simple present (regular verb)

present progressive

Bitte üben! ☆☆☆

present progressive (present progressive signal)

Hervorragend. Weiter so! ☆☆☆

Figure 2: Learner model

errors, we know which form served as the basis for generating an erroneous form, thus we know the linguistic nature of the correct target form for each error and can link an error to a specific target concept. With this graph, students at a glance can identify their diagnosed strengths and weaknesses for a target construct.

If they want to receive more detailed information, they can scroll down and for each of the specific topics displayed in the wind rose chart, they can request more details. As can be seen in the lower half of Figure 2, for each competence displayed in the wind rose chart, there is a node that the learner can unfold. In addition to a more fine-grained coloring (green, yellow, red), this competence indicator also shows a 3-way star rating indicating the proportion of uptake that resulted in the correction of a diagnosed error, compared to those cases where a student did not show successful uptake for this construct. Uptake in this scenario is defined according to Lyster and Ranta ([10]): if the learner in an attempted repair manages to correct the error targeted by the feedback message, this interaction is counted as successful uptake, independent of whether the answer still contains other errors. It thus serves as a measure of how well this student is able to make use of specific feedback regarding this target construct.

The focus of the screen shot in Figure 3 illustrates further information that can be requested for each construct. In the top half, students can see how their correct usage and errors developed over time, with each value on the x-axis corresponding to one day since the first time they used the specific construct. In this display, the colors indicate dif-

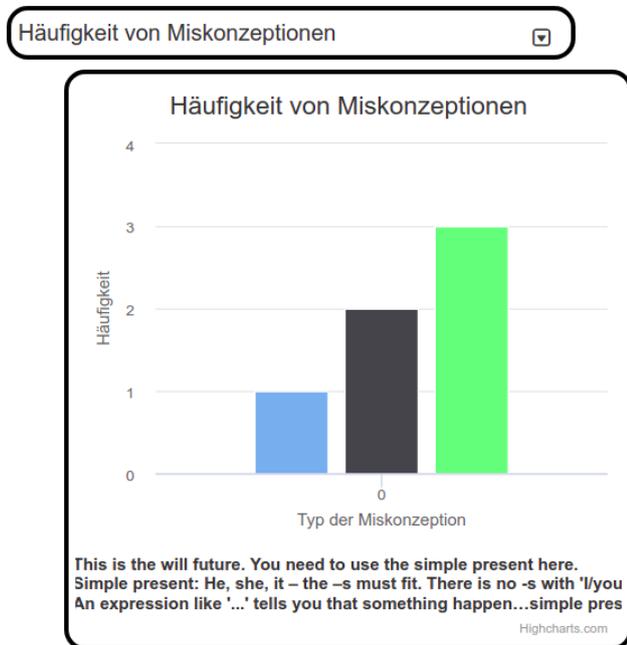
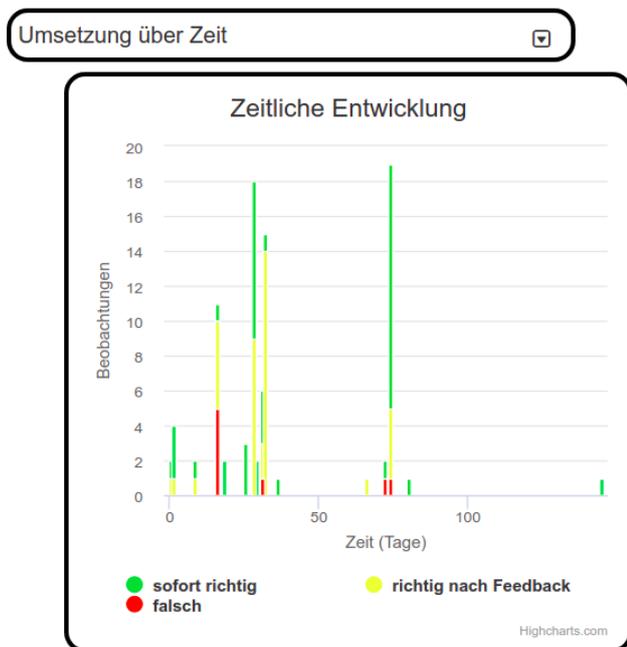


Figure 3: Temporal development and frequency of misconceptions in learner model

ferent uptake metrics. Green here indicates attempts for this target construct where the correct answer was given at first try. Yellow represents interaction sequences that ultimately resulted in the correct answer but required feedback to reach this state. Red data points are observations where the learner did not reach the correct answer for prompts expecting this target construct to be realized.

In the lower part of Figure 3, students see a bar chart that informs them about the frequency of the misconceptions they exhibited for the target construct. For one grammar con-

struct to be acquired, it is possible to make different errors, i.e., linguistically different ways of getting the target form wrong. Students with this perspective have the ability to see their typical/frequent errors for a specific grammar topic.

The learner model that is accessible through the FeedBook interface also supports construct-specific, proficiency-dependent sequencing of exercises. As discussed above in the context of Figure 2, each specific grammar concept is visualized as a node that is colored using the usual traffic light color system. The user can unfold the information on this concept further by clicking on the function “suggestions for exercises” below the concept node. In Figure 4, the student did this for the *simple present* concept, which was shown in green,

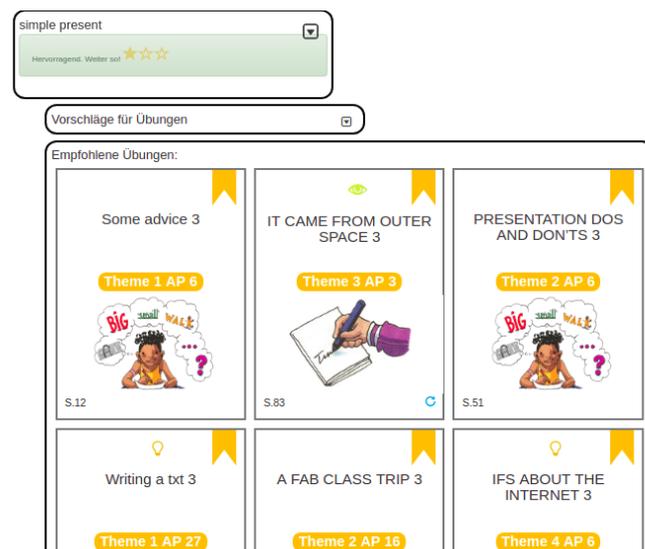


Figure 4: Proficiency-appropriate sequencing of material in learner model

showing already good mastery. The system now proposes exercises to the student that target the specific linguistic construct the student selected and are appropriate to the student’s current level of proficiency as modelled by the system. In the concrete example, the learner has shown high proficiency in using the simple present, therefore the system suggests exercises at the highest difficulty level (as indicated by the suffix “3” in the exercise titles). Where a learner has demonstrated medium competence (yellow node), the system suggests exercises of level 2, and exercises at level 1 for low competence (red node).

The goal of this approach is to present students with exercises in their Zone of Proximal Development ([23]). The system offers 158 exercises at three difficulty levels each that form the basis for this system-recommended automatic sequencing. The classification of the exercises in the three difficulty levels is taken from the publisher’s materials. With the help of EDM, based on the performance data from the study, we plan to empirically validate the externally assigned labels of difficulty and revise them were necessary. In the future, this should also make it possible to develop a more general, parametrizable model of task difficulty to support the manual or even automatic generation of tasks at different levels, in the spirit of dynamic difficulty adaptation ([15]).

In the literature, different approaches have been proposed that extend frequency-based (re)presentations of learner data ([4]). The proposals range from user stereotype models to which learners are assigned, via automatic clustering of learners into groups, fuzzy approaches modeling uncertainty, growth models weighting information differently over time, to Bayesian approaches with probabilistic dependencies. For these models the complexity lies in the data aggregation techniques employed. In the approach we are presenting here, a substantial part of the complexity of the learner model in essence is outsourced to the feedback generation method. The information grounded in the complex feedback generation process integrating language, language learning and curriculum expertise ([19, 24]) supports a comparatively simple frequency-based approach to learner modeling.

3.2 Teacher Perspective on the Class

For teachers, the most important questions in the educational context are: Where are my students in terms of proficiency? What are the constructs they are having problems with, across and in specific exercises?

While aggregated versions of the learner model described in the previous section can address the question regarding the language proficiency, the second question regarding typical errors is starting to be answered in the FeedBook using an interface that allows teachers to see how the students performed on each prompt of an exercise. In the task field performance interface, the teacher can select an exercise and click on any input field, which results in a popup window appearing, such as the one shown in Figure 5. At the top, the teacher is presented with a pie chart that lists the different types of errors were committed by their students for

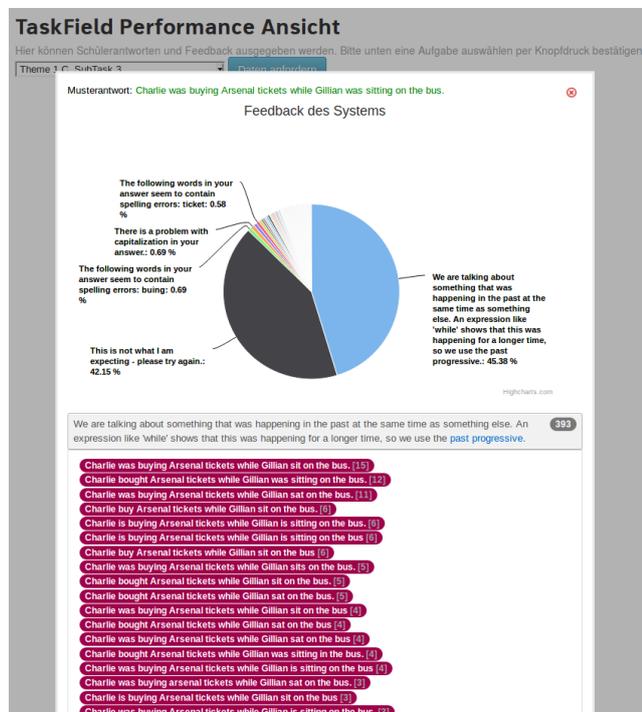


Figure 5: Prompt-specific learner answer analysis

this input field. In the concrete example, the teacher can see at a glance that the students had most difficulty with the past progressive. Below the pie chart, a teacher can then, for each slice of the pie, see the concrete number of times this error was made (here: 393 times). Upon clicking on the element, they can see the types of student answers that led to this feedback, ranked by decreasing frequency (shown in brackets). The data used as the basis for these visualizations consist of intermediate log data collected while students were working on the exercise. Note that without this interface teachers would be limited in their insights into misconceptions of their students since this interface uses the learning process data, whereas in the submitted exercises the teachers only see the final answer the student submitted after having reacted to the system feedback. Without having to grade homework themselves, this perspective thus allows the teacher to obtain an informed overview of the typical problems student faced while working on the homework, which can serve as basis for discussion in the next class.

3.3 Material Designer Perspective on Tasks

For material designers, the most important questions in the educational context are: Do the learners work with an exercise? Is the exercise designed in a way that learners actually benefit from doing the exercise? Is an exercise understandable, too easy, or too hard?

In order to address these questions, we provide perspectives for material designers in FeedBook. For each exercise, material designers can view different metrics aggregating the performance of all users in the system. Figure 6 illustrates the functionality. In the graph at the top, material designers can see the number of interactions on the x-axis against

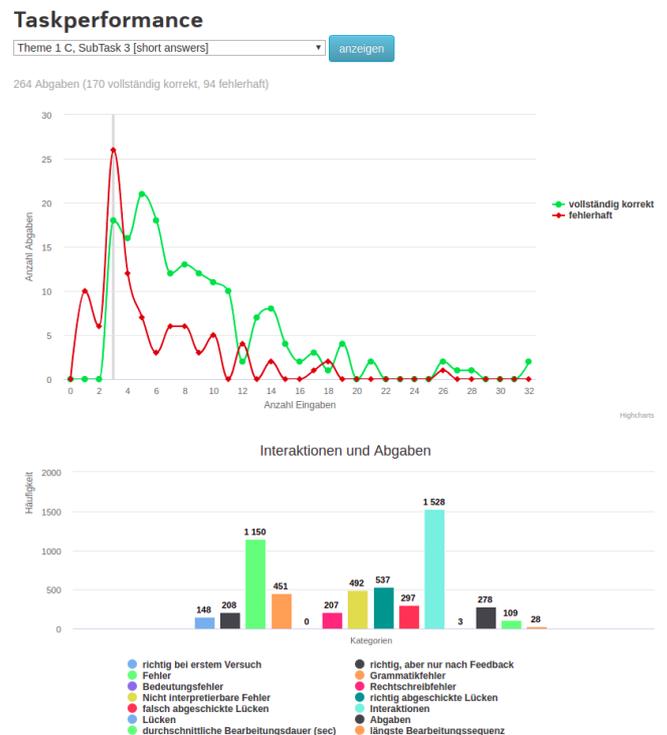


Figure 6: Interaction performance metrics

feature	description	interpretation
correct at first try	number of instances students filled in a correct answer at their first attempt	previous knowledge
errors	number of diagnosed errors	misconceptions
incorrect answers	number of answers submitted to teacher with errors	cases where the system was unable to lead to a correct solution
correct answers	number of answers submitted to teacher without errors	previous knowledge or successful uptake
interactions	total number of interactions with feedback mechanism	number of times learners requested feedback
submissions	number of submission to teachers by students	size of data set underlying visualization
gaps	number of gaps filled out	coverage of learner submission for task
average time	average time taken by learners for exercise	time on task
correct only after feedback	number of prompts where feedback led to correct solution (uptake)	number of cases where system feedback was effective
default feedback	learner answers with no associated diagnosis	lack of coverage of specific system feedback
meaning errors	diagnosed missing or additional information	semantic misunderstandings
grammar errors	diagnosed grammatical misconceptions	morphological or syntactic misunderstandings
spelling errors	diagnosed misspellings	orthographic misunderstandings
longest sequence	maximum number of steps for prompt	perseverance of learners for reaching correct answers

Table 1: Uptake metrics presented to material designers

the number of submissions on the y-axis. Each data point thus represents how many students submitted an exercise after that many interactions. Furthermore, each point is either green or red. Green indicates that the submission only contained correct answers, whereas red means that one or more learner answers contained a form that was not identical to a target answer. The vertical grey bar in the figure indicates the number of answers required to complete this exercise. To the left of the grey bar, there are only red data points since it is not possible to submit an exercise without having filled all answer slots. To the right of the grey line, the graph for this exercise shows the green line generally above the red line. This indicates that this exercise is well-designed for this learner population in that learners do not get it right immediately, but after interacting with the system more learners succeed in submitting correct answers than incorrect ones.

Below the interaction performance line chart in Figure 6, material designers can see a bar chart displaying uptake analytics. The diagram visualizes the metrics characterized in Table 1 for the selected exercise. These metrics describe the proportion of different types of errors (grammar/meaning/spelling/other) as well as the effectiveness of feedback as reflected in uptake (i.e., whether learners were able to reach a correct answers after feedback). Other measures provide insights into learner motivation, as reflected in the longest interaction sequence or the average time on task. As a next step, we plan to use such data as input to educational data mining algorithms clustering exercises by learner performance, and to cluster learners into learner groups based on the profiles extracted for individual learners.

In addition to the performance line chart and the uptake analytics, material designers have access to a pie chart with typical errors (cf. Figure 5), but with data from all users (in

anonymous form), rather than those in a specific class.

4. CONCLUSIONS AND FUTURE WORK

We illustrated how different perspectives on learner and interaction data made accessible by different user interfaces can help address demands of students (with a learner model), teachers (with a view of typical errors), and material designers (with a view on general user performance). The LA functions provide insights into the learning process based on interaction data representing incremental learner steps. Without these functions, learners, teachers and material designers would lack important information that is not available when only the student submissions as final product are available.

Future work planned in terms of system development includes implementing improved proficiency-dependent sequencing of materials taking into account the learner’s competence for all grammar constructs rather than just the currently focused on. We also plan to improve the interface side, where instead of requiring learners to navigate through different grammar constructs, the system could automatically combine the evidence and simply provide a user interface element “next exercise”. We also plan to generally improve the accessibility of the learner model. We are also considering extending the learner model in the directions argued for in [1]. This would allow us to not only model language misconceptions but also to what extent a learner is capable of completing tasks requiring specific strategies. Furthermore, we look forward to continuing the discussion with the teachers using the system so that the learning analytics functions address actual real-life needs. One envisaged extension is an aggregated perspective on the learner models of all students in class to diagnose trends in their competence and misconceptions across tasks.

5. REFERENCES

- [1] L. Amaral and D. Meurers. 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, 2008.
- [2] S. Bull and J. Kay. *Student models that invite the learner in: The SMILI open learner modelling framework*. Citeseer, 2006.
- [3] S. Bull and T. Nghiem. Helping learners to understand themselves with a learner model open to students, peers and instructors. In *Proceedings of workshop on individual and group modelling methods that help learners understand themselves, International Conference on Intelligent Tutoring Systems*, volume 2002, pages 5–13. Citeseer, 2002.
- [4] K. Chrysafiadi and M. Virvou. Student modeling approaches: A literature review for the last decade. *Expert Systems with Applications*, 40(11):4715–4729, 2013.
- [5] R. Ellis. *Task-based Language Learning and Teaching*. Oxford University Press, Oxford, UK, 2003.
- [6] R. Ferguson. Learning analytics: drivers, developments and challenges. *International Journal of Technology Enhanced Learning*, 4(5/6):304–317, 2012.
- [7] T. Heift. Corrective feedback and learner uptake in CALL. *ReCALL*, 16(2):416–431, 2004.
- [8] M. H. Long. The role of linguistic environment in second language acquisition. In W. C. Ritchie and T. K. Bhatia, editors, *Handbook of second language acquisition*, pages 413–468. Academic Press, New York, 1996.
- [9] C. Lynch, K. Ashley, V. Alevan, and N. Pinkwart. Defining ill-defined domains; a literature survey. In *Proceedings of the workshop on intelligent tutoring systems for ill-defined domains, held during the 8th international conference on intelligent tutoring systems*, pages 1–10, Jhongli, Taiwan, 2006.
- [10] R. Lyster and L. Ranta. Corrective feedback and learner uptake: Negotiation of form in communicative classrooms. *Studies in Second Language Acquisition*, 19 n1:37–66, 1997.
- [11] D. Meurers. Natural language processing and language learning. In C. A. Chapelle, editor, *Encyclopedia of Applied Linguistics*, pages 4193–4205. Wiley, Oxford, 2012.
- [12] D. Meurers, K. De Kuthy, F. Nuxoll, B. Rudzewitz, and R. Ziai. Scaling up intervention studies to investigate real-life foreign language learning in school. *Annual Review of Applied Linguistics*, 39, 2019. To appear.
- [13] D. Meurers and M. Dickinson. Evidence and interpretation in language learning research: Opportunities for collaboration with computational linguistics. *Language Learning*, 67(2), 2017.
- [14] D. Meurers, K. D. Kuthy, V. Möller, F. Nuxoll, B. Rudzewitz, and R. Ziai. Digitale Differenzierung benötigt Informationen zu Sprache, Aufgabe und Lerner. Zur Generierung von individuellem Feedback in einem interaktiven Arbeitsheft. *FLuL – Fremdsprachen Lehren und Lernen*, 47(2):64–82, 2018.
- [15] I. Pandarova, T. Schmidt, J. Hartig, A. Boubekki, R. D. Jones, and U. Brefeld. Predicting the difficulty of exercise items for dynamic difficulty adaptation in adaptive language tutoring. *International Journal of Artificial Intelligence in Education*, 2019.
- [16] K. Petersen. *Implicit Corrective Feedback in Computer-Guided Interaction: Does Mode Matter?* PhD thesis, Georgetown University, 2010.
- [17] C. Romero and S. Ventura. Educational data mining: A survey from 1995 to 2005. *Expert systems with applications*, 33(1):135–146, 2007.
- [18] B. Rudzewitz, R. Ziai, K. De Kuthy, and D. Meurers. 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*, 2017.
- [19] B. Rudzewitz, R. Ziai, K. De Kuthy, V. Möller, F. Nuxoll, and D. Meurers. Generating feedback for English foreign language exercises. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*, pages 127–136. ACL, 2018.
- [20] K. E. Sabo, R. K. Atkinson, A. L. Barrus, S. S. Joseph, and R. S. Perez. Searching for the two sigma advantage: Evaluating algebra intelligent tutors. *Computers in Human Behavior*, 29(4):1833–1840, 2013.
- [21] G. Siemens and R. S. Baker. Learning analytics and educational data mining: towards communication and collaboration. In *Proceedings of the 2nd international conference on learning analytics and knowledge*, pages 252–254. ACM, 2012.
- [22] V. Slavuj, A. Meštrović, and B. Kovačić. Adaptivity in educational systems for language learning: a review. *Computer Assisted Language Learning*, 30(1-2):64–90, 2017.
- [23] L. S. Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard University Press, Cambridge, MA, 1978.
- [24] R. Ziai, B. Rudzewitz, K. De Kuthy, F. Nuxoll, and D. Meurers. Feedback strategies for form and meaning in a real-life language tutoring system. In *Proceedings of the 7th Workshop on Natural Language Processing for Computer-Assisted Language Learning (NLP4CALL)*, pages 91–98. ACL, 2018.