VU Research Portal

A Robot Math Tutor that Gives Feedback

Hindriks, Koen V.; Liebens, Sander

published in
Social Robotics
2019

DOI (link to publisher)
10.1007/978-3-030-35888-4_56

document version
Publisher's PDF, also known as Version of record

document license
Article 25fa Dutch Copyright Act

Link to publication in VU Research Portal

citation for published version (APA)
Hindriks, K. V., & Liebens, S. (2019). A Robot Math Tutor that Gives Feedback. In M. A. Salichs, S. S. Ge, E. I. Barakova, J.-J. Cabibihan, A. R. Wagner, Á. Castro-González, & H. He (Eds.), Social Robotics: 11th International Conference, ICSR 2019, Madrid, Spain, November 26–29, 2019, Proceedings (pp. 601-610). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 11876 LNAI). Springer. https://doi.org/10.1007/978-3-030-35888-4_56

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:
vuresearchportal.ub@vu.nl

Download date: 29. Mar. 2024

A Robot Math Tutor that Gives Feedback

Koen V. Hindriks[1(✉)] and Sander Liebens[2]

[1] Faculty of Science, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
k.v.hindriks@vu.nl
[2] Department of Computing Science, Delft, The Netherlands
sander.liebens@gmail.com

Abstract. We report on the exploratory design and study of a robot math tutor that can provide feedback on specific errors made by children solving basic addition and subtraction problems up to 100. We discuss two interaction design patterns, one for speech recognition of answers when children think aloud, and one for providing error-specific feedback. We evaluate our design patterns and whether our feedback mechanism motivates children and improves their performance at primary schools with children ($N = 41$) aged 7–9. We did not find any motivational or learning effects of our feedback mechanism but lessons learnt include that the robot can execute our interaction design patterns autonomously, and advanced algorithms for error classification and adaptation to children's performance levels in our feedback mechanism are needed.

1 Introduction

Feedback is one of the most powerful influences on learning and achievement [4]. In this paper we present the design of and evaluate a social robot that can provide error-specific feedback on answers to math problems. We focus on the domain of addition and subtraction, for sums and differences up to 100. Children aged 7–8 are being introduced to these problems with the aim to be able to fluently add and subtract two-digit numbers. We deploy a social robot to provide additional training to children aged 8–9 at school to help children to learn to automate the strategies they use for adding and subtracting numbers.

Social robot tutors have recently become more popular in elementary schools, for example, to teach language and mathematics skills. Children are motivated to work with social robots. Their physical presence compares favourably with other tutoring technologies (e.g., using games on tablets) as it appears to increase cognitive and affective outcomes [2,10]. We briefly discuss a few related studies that use social robots to teach mathematics. [5] combines a variety of 30 different arithmetic tasks with an engaging game in which children aged 9–10 have to imitate the gestures a Nao robot performs and finds evidence that suggests that challenging children more helps them reach higher problem levels. [3] studies the effects of a Darwin robot as a tutor for children aged 13–18 that were asked to complete algebra problems. They find that the use of verbal engagement (instead of only nonverbal cues) enhances test performance most on these problems. [7]

© Springer Nature Switzerland AG 2019
M. A. Salichs et al. (Eds.): ICSR 2019, LNAI 11876, pp. 601–610, 2019.
https://doi.org/10.1007/978-3-030-35888-4_56

compares the effects of a Nao robot executing a social versus an a-social tutoring strategy with children learning about prime numbers. Their results suggest that a physical robot leads to improved learning but also may lead a child to pay more attention to the robot's social behaviour than the lesson content. In these studies, no effort is made to identify the type of error that is made and no detailed feedback on how to correct that error is provided so the child can learn how to avoid it next time, which is the focus of the work reported on here.

Feedback is information about how you are doing in your efforts to reach a goal [4,11,17]. We build on the conceptual model proposed in [4] to design a feedback mechanism that the robot can use to help a child understand what kind of error it made (if it did) in how it produced an answer to a math problem. According to [4], feedback directed to the self is the least effective, while self regulated feedback and feedback about the process of a task are most powerful in terms of mastery of tasks. Immediate feedback, moreover, gives better results than feedback that is delayed [14,17].

The main idea we explore in this paper is that the robot not only automatically detects a mistake but also identifies what kind of mistake has been made and explains how to avoid this mistake next time. The robot uses speech to interact with the children and we present an interaction design pattern to handle the typical thinking aloud behaviour of these children. We also design a feedback mechanism and compare task performance, affection, and the interaction with the robot of a control group that did not receive feedback with a feedback group that received explicit feedback on mistakes made from the robot. We hypothesise that (i) the score on math problems is higher for children that receive feedback from the robot, and that (ii) children that receive feedback like the robot more.

The paper is organised as follows. Section 2 introduces our robot interaction design and Sect. 3 our feedback mechanism. Sections 4, 5, and 6, present method, results, and discussion of our study. Section 7 concludes the paper.

2 Interaction Design

We designed a robot math tutor called Pixel that is able to teach children how to perform the basic operations addition and subtraction up to 100.[1] At the start of a session, Pixel introduces itself and welcomes the child, after which it explains that it is going to pose some addition and subtraction problems which the child has already been practising with for some time in the classroom. This strategy of explaining a subject before engaging in it has been shown to be an effective teaching strategy [12,15]. The robot has been programmed to select problems from seven different categories (cycling through these categories).[2]

[1] This is task CCSS.MATH.CONTENT.2.NBT.B.5 in the Common Core Standard for Grade 2 (age 7–8) http://www.corestandards.org/Math/Content/2/NBT/B/5/.

[2] The following 7 categories are used: "Passing tens" (e.g., $7 + 6$), "Adding tens and units" (e.g., $37+31$), "Adding tens and units, passing tens" (e.g., $67+14$), "Through tens" (e.g., $12 - 5$), "Remove tens and add later" (e.g., $53 - 2$), "Tens minus tens and units" (e.g., $38 - 17$), and "Tens and units, through tens" (e.g., $46 - 18$).

Providing an Answer. During a pilot, we found that children may engage in thinking aloud while trying to compute an answer. This poses a problem for the robot when it starts listening for an answer. It is hard for the robot to identify exactly which of the numbers a child mentions is supposed to be the answer. To deal with this issue, we developed an interaction design pattern (see Table 1) called *Touch-Based Speech Activation* for providing an answer to the robot [6].

Table 1. Interaction design pattern: Touch-Based Speech Activation

Problem	When asked to answer a question by a robot, children may engage in thinking aloud while trying to compute the answer to the question. Children's speech while thinking aloud is harder to recognise as speech volume, for example, is varied more. Both the longer and more complicated speech produced (instead of providing only the answer) and the variation in speech parameters complicates the natural language understanding, in particular the identification of the answer
Principle	We do not want to restrict children in the way they compute an answer, and allow them to engage in thinking aloud and other interaction (e.g. asking another child sitting next to them). Instead, to provide an answer, a child is asked to indicate it is ready and focused to provide an answer.
Solution	A child is asked to indicate that it thinks it knows the answer to a question by means of a touch sensor. Touching the sensor will activate speech recognition and the robot will then listen for an answer for a specified period of time

We used the feet bumper sensors of the Nao robot to activate speech recognition and used a window of 3 seconds to listen for an answer. Note that the interaction pattern is quite generic and can be used for any type of question, and only requires a robot with a touch sensor. Because we did not use any visual support e.g. using a tablet for displaying the question and children sometimes fail to understand or memorise the question (or simply want to be reassured they answer the right question), we found it also useful to provide them with a touch mechanism for asking the robot to repeat a question. We enabled children to ask the robot to repeat a question by pressing buttons on the robot's head.

Rewards. To motivate children, rewards are provided if the correct answer is given. Rewards have been implemented by randomly having the robot say one of four compliments while displaying a rotating rainbow colour pattern using the LEDs in the robot's eyes. The four texts we used are "Great job! That indeed is the correct answer!", "Correct again! You're doing great!", "You're the best! That is the correct answer!", and "Yes, correct! Let's go for the next one!". The rewards are not part of the feedback mechanism discussed below and used to motivate all children (whether or not they receive extensive feedback or not from the robot). At the end of a session, the robot also informs the child how many exercises were answered correctly in the first try, thanks the child for paying attention and completing the exercises with the robot, and, as a final reward, first teaches the child how to perform the short *Chu Chu Wa* dance and then invites the child to perform it together.[3]

3 Feedback Targeting the Type of Mistake Made

One of the simplest feedback mechanisms only informs the child whether the answer it provided was right or wrong. We used this mechanism for children in the control group in our study to provide feedback on each answer. After providing an incorrect answer in a first attempt, children always were invited to try a second time. If in a second attempt a child also failed to provide the right answer, the robot would indicate this and provide the right answer for the current problem. In case of a good answer a reward (see above) is provided.

A second, more elaborate mechanism that we used with the feedback (test) group in our study is based on the model of [4]. The mechanism we propose is based on the "How are we going?" question and aims to provide feedback at the process level of this model. We have made these choices as feedback is effective when it consists of information about progress, and/or about how to proceed [4]. Speech is used to provide verbal feedback, inspired by [15] who argue that a dialog in teaching is more effective. The main aim of our feedback mechanism is to help children understand *what* they did wrong by providing relevant feedback on the specific error that most likely was made. See the algorithm discussed below for detecting the type of error made for details. If the kind of error made is identified by this algorithm, feedback is adapted specifically to this error to explain what went wrong and to help the child understand how it can correct its mistake. If the algorithm cannot classify the error type, the robot will indicate the answer is wrong the first attempt but after a second attempt will explain how to compute the correct answer using the jump strategy [1].

Detecting Errors in Answers to Math Problems. In order to give specific feedback, the robot should be able to diagnose what kind of error is made when a child provides an incorrect answer. The literature provides some insight into the types of errors that are made when children use the jump and/or split strategy [1], but does not detail how to mechanically detect such errors. We therefore have taken a pragmatic approach and implemented an algorithm to detect some of the errors in the addition or subtraction domain that are easy to detect. This would suit the purpose of evaluating our feedback mechanism assuming these errors occur frequently enough. We have included the following four commonly made errors: (i) "missing the last step", e.g., $7+5 = 7+3 = 10$, forgetting to add 2, (ii) "visualising a number incorrectly", e.g., switching numbers and use 83 instead of 38, (iii) the "split and add error", e.g. $32 - 14 = 22$, by first subtracting 10 from 30, then subtracting 2 from 4, because the other way around seems impossible, and (iv) "using the wrong operator", e.g., $7 + 5 = 7 - 5 = 2$.

[3] See https://www.youtube.com/watch?v=3T5cPaZIW8M for a video of the dance.

Table 2. Interaction Design Pattern: Feedback Targeting Specific Error

Problem	Children make a variety of errors when answering math problems, in e.g., the addition and subtraction domain. Feedback is more effective if it can target the error made more specifically
Principle	We want to provide feedback that specifically focuses on the type of error a child makes. The robot should be able to explain *what* went wrong to help the child understand its mistake but also allow a child to correct the error to learn from it.
Solution	An algorithm is used to classify the type of error a child makes when answering math problems. Feedback is designed to specifically explain what the child did wrong to help the child understand how to fix the error in a step-by-step fashion. If the type of error cannot be classified, simply provide feedback that the answer is incorrect. In any case, allow the child to retry and provide an answer to the same problem again. If an incorrect answer is provided for the second time, indicate this and have the robot explain how the correct answer can be computed

For each of the four errors the algorithm can detect, different feedback is designed. For "missing a step", the child is told which step it may have missed. For "wrongly visualising a number", the child is told it may have switched numbers and how to fix that. For "split and add error", the child is told it may have switched units in the different numbers. And, finally, for "using the wrong operator", the child is told it may have used the wrong operator and which operator should have been used in the problem.

Pilot. A primary school teacher was asked to assess the wording used by the robot and whether children could understand the feedback. This resulted in a few changes to the wording used. A pilot with a small group of children at a day care centre confirmed that children could understand the robot's feedback.

4 Method

4.1 Design

We used a between-subject study design. The independent variable is the feedback provided (only indicating whether answer is correct or not, or also providing error-specific feedback). The dependent variables are the test results of the math problems (amount of correct answers, ranging from 0 to 20) and the PANAS scores (sum of the weights, ranging from -40 to 40). The interaction design patterns were implemented across all conditions. All the interactions that were included in the experiment were used to evaluate these patterns.

4.2 Participants

A total of 41 children aged 7–9 ($M = 8.17$), 23 male and 18 female, from two primary schools in The Netherlands participated (20 children from one school, 21 from the other). We obtained approval from the Ethics Committee and consent of the children's parents beforehand. Children were split into a feedback and control group with on average the same level of math performance using the results of a test of 20 math problems administered in class before the experiment and an average math grade provided by the teacher. The feedback group consisted of 13 boys and 7 girls and the control group of 10 boys and 11 girls.

4.3 Materials

A human-like Nao robot (V5) from SoftBank Robotics was used, with a standard laptop to control it. Three different paper-and-pencil tests with 20 math questions each were used as pre- and post-tests. A small survey of 5 open and 6 Likert-scale (5-point) questions about math and robots in general, and about their math session with the robot was used. The survey for the feedback group had two additional questions about the feedback from the robot. We used the PANAS form to measure affection [16]. Finally, an observation sheet was used for notes by the experimenter.

4.4 Procedure

A week before the individual sessions with the robot, all children took part in an hour-long class in which the robot was introduced. At the end of this class they were asked to complete a math pre-test. During the math sessions, the robot was standing on a table, within reach of the child sitting in front of it. A brief explanation of the session was given before a training session was started to make the child familiar with how to interact with and answer math problems posed by the robot. Help was provided when needed. The training session was repeated only once for a child who could not complete it without repeated assistance. After the training session, children also completed a PANAS form. The first session took about 45 min with 20 min in which the robot asked children to answer math problems. All children were told by the robot whether the answer was correct or not, but the feedback group also received detailed feedback on the type of mistake that was made, if the robot could identify it. All exercises with given answers of the first session were stored in a database. A week later children participated in a second session of about 30 min which did not include training, with again 20 min of math problems. In this session the robot selected math problems from categories in which a child made (more) mistakes than other categories during the first session. After each math session, children were asked to complete a PANAS form and a short survey. After the second math session children completed a first math post-test. Two weeks later, children completed a second post-test.

5 Results

The control group completed 638 in the first and 515 math problems in the second session, whereas the feedback group completed 477 in the first and 420 in the second session. Children in the control group completed significantly more questions ($M = 30.4$) in the first than in the second session ($M = 24.5$, $t(20) = 4,353$, $p = 0,0003$). No difference was found in the feedback group ($M = 23.9$ questions in the first and $M = 21.0$ in the second session). Using One-Way ANOVA, we find that the control group completed significantly more exercises in the first session than the feedback group ($F(1, 39) = 7,321$, $p = 0.010$) but no difference was found in the second session ($F(1, 39) = 1,679$, $p = 0.203$).

The control group gave 254 incorrect answers (including repeated tries) for 164 questions (25.7%) in the first and 261 incorrect answers for 157 questions (30.5%) in the second session. The feedback group gave 221 incorrect answers for 143 questions (30.0%) in the first and 219 incorrect answers for 139 questions (33.1%) in the second session. All mistakes that could be identified by our detection algorithm as fitting in one of the error categories we defined were identified as such and all categories occurred at least once but our algorithm was only able to recognise 15 out of 254 incorrect answers given by the feedback group in the first session and 12 out of 219 in the second session. Most of the mistakes made thus could not be classified by the algorithm, often because, for example, answers given seemed 'random' and any pattern was hard to detect. Upon the given feedback, 44.4% of the time children were able to produce a correct answer. Because feedback was also given to children who gave incorrect answers after a second try, feedback was given 158 times (35.9%) in total in both sessions.

We did not find any significantly different results on the three math tests that were taken before and after the experiment between the control and feedback group. The results of the first math test were not significantly different from the second ($p = 0.316$), and those of the second were also not significantly different from the last one ($p = 0.630$). Similarly, we used Mixed ANOVA to analyse the PANAS scores of the children, but did not find any significant differences between the feedback and the control group nor between sessions.

Children needed to get used to the interaction mechanism that we designed, although they hardly ever (5 out of 3994 times) forgot to press the feet before telling the robot their answer. The main issues were that children provided their answer too late and the speech-to-text module sometimes failed. Children provided an answer too late, more in the first than the second session (2.51 times on average versus 0.66, a significant difference, $t(49) = 2.94$, $p = 0.005$, with 3 children who had more issues in this regard but only in the first session). The robot misinterpreted in the first session on average $3, 4$ correct answers per child, significantly more often than the $1, 7$ answers that were incorrectly transcribed to text in the second session ($t(80) = 2.67$, $p = 0.009$). Children also learned to ask the robot to repeat a math problem by touching the robot's head more in the second (10.6) than the first session (6.8, $t(75) = -3.28$, $p = 0.002$).

Children indicated in the survey that they liked working with the robot and that they felt they learned something from the robot. Only in the control group, we found a significant increase in rating of the statement "I think Pixel [i.e. the robot] has taught me something", with an average of $M = 3.81$ out of 5 after the first session versus $M = 4.19$ after the second session ($t(20) = 2.36$, $p = 0.029$).

6 Discussion

Our detection algorithm was only able to classify a few of the many mistakes that children made, limiting the feedback that children received from the robot. Although theory has identified several categories of commonly made mistakes, many of the mistakes we found in practice did not fit into these categories and another approach is needed to handle these. We found evidence that suggests that feedback is more appreciated by children who have more difficulty with math problems and appreciate the fact that the robot put no time pressure on providing an answer. Children who perform well at math sometimes got impatient with the robot that explained them things they already know.

We could not establish any learning effects by means of the math tests children completed before and after the experiment. We thus found no evidence to support our hypothesis that feedback provided by a robot improves performance on math problems. This may have been partly due to the high scores on the pre-test completed before the experiment leaving little room for improvement in the other tests but equally likely is that a few sessions with a robot are insufficient to find any learning effects on math questions. Although most children indicated that the robot improved their math skills and that they learnt from the robot, a few said that the robot did not teach them anything new. Some of the better performing children also said it took the robot too long to provide feedback, and they would have preferred to be able to move on to the next question quicker.

The PANAS results also do not support our hypothesis that a robot that provides more feedback is liked more. Interestingly, however, some children said they liked working with the robot better than doing math on a tablet because the robot did not use a timer as the tablet apps do. The tablet apps for math force children to produce answers quickly while the robot gave children as much time as they needed. Teachers at the schools that participated also were very enthusiastic about the use of a robot math tutor.

Children need time to learn how to interact with a robot, which became clear from the two sessions in which children's interaction with the robot significantly improved: in the second session children were quicker to provide the robot with an answer, they asked the robot more often to repeat the question by touching its head, and they learnt how to talk to the robot when providing their answers to it. Some of the students were shy and had trouble communicating with the robot at first, but this improved in the second session. The robot's misunderstandings were frustrating, and children indicated after the sessions that the robot should listen better to their answers. Issues with speech recognition remain a big issue in child-robot interaction [8]. Overall, however, we can say that the combination of speech and touch created a robust interaction mechanism as all children were able to successfully complete the sessions without any intervention from the researcher who was present.

Limitations. Our results indicate that a few sessions over a two week period are insufficient to measure any learning effects of the robot on the children's performance on math problems. It seems hard to ascribe any effects specifically to the robot when the teacher spends much more time teaching the children.

7 Conclusion

Our approach to identify specific errors in addition and subtraction problems, as was to be expected, has turned out to be somewhat naive. We have overestimated the frequency of occurrence of the types of problems our algorithm is able to detect. Moreover, our method does not detect the strategy used by a child to solve the problem. Even though it is not clear from the literature whether feedback is best given using the same strategy as the child uses, we are also unable to do this. One important lesson of our work therefore is that there is a need for more detailed knowledge about which errors children make in the addition and subtraction (and other) domain(s) and how we can detect these and the strategies children use in this domain. As [1] indicates, there is still much to learn about this domain and this is still the case even though some progress has been made more recently in the automated diagnosis of subtraction bugs [9].

It is hard to assess the effectiveness of our feedback approach given the low frequency of addition and subtraction errors that the robot was able to detect. We did not find any evidence to support our hypothesis that feedback improves children's performance in the addition and subtraction domain. One issue, unrelated to the robot here, is that the children in our study already did quite well on the problems we presented. We also did not find any evidence that feedback increases affection towards the robot math tutor. We did receive very positive comments from children on learning with the robot though.

The interaction mechanism was robust and the robot was able to execute our interaction design patterns autonomously. Children were always able to complete a session without any need for human intervention. We found that the interaction with the robot in the second session improved compared to the first. With regards to our first Touch-Based Speech Activation pattern, fewer children were late in providing an answer and the robot made fewer speech recognition mistakes. Children needed some time but overall were able to quickly learn how to interact with the robot. With regards to our second Feedback pattern, it appears that differences in children's performance also should be reflected more in the feedback mechanism. Children that perform better, for example, become impatient when they have to listen to the rather long feedback of the robot in case of an error and should be able to proceed quicker to the next problem. It would also be interesting to adapt the problem level to the performance level of a child [13].

Finally, we have focused on verbal feedback and only used non-verbal rewards (LEDs, dance) but non-verbal behaviour can also be used to give feedback. The integration of such behaviour, e.g., using nodding or head shaking to steer a child in the right direction [12], may be an interesting direction for future work.

References

1. Beishuizen, M.: Mental strategies and materials or models for addition and subtraction up to 100 in Dutch second grades. J. Res. Math. Educ. **24**(4), 294–323 (1993)
2. Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., Tanaka, F.: Social robots for education: a review. Sci. Roboti. **3**(21), eaat5954 (2018)
3. Brown, L., Howard, A.M.: Engaging children in math education using a socially interactive humanoid robot. In: 2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids), pp. 183–188, October 2013
4. Hattie, J., Timperley, H.: Review of educational the power of feedback. Rev. Educ. Res. **77**(1), 81–112 (2007)
5. Janssen, J.B., van der Wal, C.C., Neerincx, M.A., Looije, R.: Motivating children to learn arithmetic with an adaptive robot game. In: Mutlu, B., Bartneck, C., Ham, J., Evers, V., Kanda, T. (eds.) ICSR 2011. LNCS (LNAI), vol. 7072, pp. 153–162. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25504-5_16
6. Kahn, P.H., et al..: Design patterns for sociality in human-robot interaction. In: Proceedings of the International Conference on Human Robot Interaction, pp. 97–104. ACM (2008)
7. Kennedy, J., Baxter, P., Belpaeme, T.: The robot who tried too hard: social behaviour of a robot tutor can negatively affect child learning. In: Proceedings of the International Conference on Human-Robot Interaction, vol. 801, pp. 67–74 (2015)
8. Kennedy, J., et al.: Child speech recognition in human-robot interaction: evaluations and recommendations. In: Proceedings of the International Conference on Human-Robot Interaction, pp. 82–90. ACM (2017)
9. Lee, J., Corter, J.E.: Diagnosis of subtraction bugs using Bayesian networks. Appl. Psychol. Meas. **35**(1), 27–47 (2011)
10. Mubin, O., Stevens, C.J., Shahid, S., Al Mahmud, A., Dong, J.J.: A review of the applicability of robots in education. J. Tech. Educ. Learn. **1**(209–0015), 13 (2013)
11. Norcini, J.: The power of feedback. Med. Educ. **44**(1), 16–17 (2010)
12. Saerbeck, M., Schut, T., Bartneck, C., Janse, M.D.: Expressive robots in education: varying the degree of social supportive behavior of a robotic tutor. In: Proceedings on Human Factors in Computing Systems (SIGCHI), pp. 1613–1622. ACM (2010)
13. Schadenberg, B.R., Neerincx, M.A., Cnossen, F., Looije, R.: Personalising game difficulty to keep children motivated to play with a social robot: a Bayesian approach. Cogn. Syst. Res. **43**, 222–231 (2017)
14. Singh, R., et al.: Feedback during web-based homework: the role of hints. In: Biswas, G., Bull, S., Kay, J., Mitrovic, A. (eds.) AIED 2011. LNCS (LNAI), vol. 6738, pp. 328–336. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21869-9_43
15. Tiberius, R.G., Billson, J.M.: The social context of teaching and learning. New Dir. Teach. Learn. **1991**(45), 67–86 (1991)
16. Watson, D., Clark, L.A., Tellegen, A.: Development and validation of brief measures of positive and negative affect: the PANAS scales. J. Pers. Soc. Psychol. **54**(6), 1063 (1988)
17. Wiggins, G.: 7 Keys to Effective Feedback. Educ. leadersh. **70**(1), 10–16 (2012)