Explaining Recommendations in E-Learning: Effects on Adolescents' Initial Trust

JEROEN OOGE*, KU Leuven, Belgium
SHOTALLO KATO*, KU Leuven, Belgium
KATRIEN VERBERT, KU Leuven, Belgium

Recommender systems are increasingly supporting explanations to increase trust in their recommendations. However, studies on explaining recommendations typically target adults in low-risk e-commerce or media contexts, and using explanations in e-learning has received little research attention. To address these limits, we investigated how explanations affect adolescents' trust in an exercise recommender on a mathematical e-learning platform. In a randomized controlled experiment with 37 adolescents, we compared real explanations with placebo and no explanations. Our results show that explanations can significantly increase initial trust when measured as a multidimensional construct of competence, benevolence, integrity, intention to return, and perceived transparency. Yet, as not all adolescents in our study attached equal importance to explanations, it remains important to tailor them. To study the impact of tailored explanations, we advise researchers to include placebo baselines in their studies as they may give more insights into how much transparency people actually need, compared to no-explanation baselines.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

Additional Key Words and Phrases: e-learning, explainable recommender systems, trust

ACM Reference Format:
Jeroen Ooge, Shotallo Kato, and Katrien Verbert. 2018. Explaining Recommendations in E-Learning: Effects on Adolescents' Initial Trust. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

People are increasingly relying on recommender systems that suggest relevant items, for example movies and music, tailored to their needs and interests. However, people are often left in the dark when it comes to why something has been recommended. In the scope of *explainable artificial intelligence* (XAI), many researchers agree that accompanying recommendations with explanations is often desirable because it can, for example, increase appropriate trust in the recommender [4, 41, 51], which in turn can increase people's willingness to adopt technologies and their outcomes [7]. Therefore, XAI and trust have become important topics in human-computer interaction.

However, the degree to which results of previous research on explaining recommender systems can be generalized is limited because of three reasons. First, studies are mostly framed in application contexts like media recommending [e.g., 8, 21, 40, 52] and e-commerce recommending [e.g., 7, 24, 48]. Other contexts, for example education [6], are underexplored. Second, most study participants are university students or young adults, resulting in scarce results for

*Both authors contributed equally to the paper adolescents (ages 11–19 [19]). Third, on a methodological level, most XAI research measures the effect of explanations by comparing recommender systems with and without explanations. However, this comparison could be unfair as recent studies suggest that the mere presence of *placebo explanations* (i.e., explanations without any meaningful content) can already increase someone's trust in a platform [17, 41].

To mitigate these limitations, we investigated how explanations affect adolescents' trust in a mathematical e-learning platform that recommends exercises. We further included placebo explanations as an extra baseline. In particular, our research questions were as follows:

- **RQ1.** Can explanations increase adolescents' initial trust in an e-learning platform that recommends exercises?
- **RQ2.** How can placebo explanations influence adolescents' initial trust in such an e-learning platform?

Our research contribution is threefold. First, we show that explaining recommendations can significantly increase initial trust in an e-learning platform if trust is measured multidimensionally. However, when measuring trust one-dimensionally, the increase is not significant, which suggests that mainly dynamically learned factors grow initial trust. Second, by comparing our explanation interface with a placebo baseline, we reveal that not all adolescents have equal needs for transparency, so tailoring explanations is essential. Third, we present a unique data set of adolescents' trust in an e-learning context, which we share publicly in the spirit of open science[1]. In sum, we hope our work inspires other researchers to target adolescents more often and study how tailored explanations impact their trust in e-learning recommender systems.

2 BACKGROUND AND RELATED WORK

This section discusses some challenges of explaining artificial intelligence, and particularly recommender systems. Then, it zooms in on trust in automated systems and previous studies about the trust effects of explaining recommendations.

2.1 Explainable Artificial Intelligence

Artificial intelligence has become ubiquitous in our time. However, ever since its resurgence, there has been a call for more algorithmic transparency because sophisticated algorithms are often 'black-boxes': the precise way in which they process vast amounts of input data to obtain an output is unclear. Not explain algorithms' outcomes may suffice for low-stakes applications like movie recommendation but becomes unacceptable in high-stakes contexts like healthcare and e-learning. *Explainable artificial intelligence* (XAI) is an umbrella term for techniques that try to explain the logic behind algorithmic decision-making, such that people can understand it, grow appropriate trust in the algorithm, and detect potential biases [27]. A substantial challenge is that XAI encompasses many intertwined topics including trust, fairness, bias, causality, accountability, privacy, and reasoning [3]. As a consequence, XAI does not yet have a widely accepted definition [16, 22, 35]. Furthermore, design requirements for explanations depend on the target audience [8, 40, 41] and context [55], and explanations can be evaluated according to several metrics [41]. Nevertheless, there is a vast literature on explanations for different algorithmic families [4, 5, 26].

2.2 Explaining Recommendations

Current XAI research often builds upon earlier research with recommender systems. For example, Herlocker et al. [29] compared different explanation designs for collaborative filtering recommenders, and showed that they could increase acceptance of recommendations. In general, explanations for recommendations come in three representational

[1] [LINK TO GITHUB]

forms [45]. (1) *Textual explanations* use natural-language phrases. Many commercial applications already employ these kinds of explanations, following patterns like "*People who liked X also liked Y*" for collaborative filtering recommenders, and "*You will like X because it has Y and Z*" for content-based recommenders. (2) *Visual explanations* use (interactive) visualizations to efficiently convey a lot of information. For example, Herlocker et al. [29] used a histogram to show how neighboring users rated a recommended movie; Tsai and Brusilovsky [53] used among others radar charts and Venn diagrams to explain similarity-based recommenders; and Bostandjiev et al. [9] visualized a music recommending process with an interactive pathway chart. (3) *Hybrid explanations* leverage both textual and visual information. For example, Gedikli et al. [21] used tag clouds in which word size encodes relevance, and Szymanski et al. [50] combined a partial dependence plot with text on how to interpret the visual information.

Overall, the recommendation algorithm usually constrains the explanation type [51]. For example, collaborative filtering recommendations cannot be explained by their inherent features. Furthermore, several trade-offs have to be made when designing explanations: Tintarev and Mastoff [51] discussed several of them and outlined seven goals for explanations, including the increase of trust.

2.3 Trust in Automated Systems

Trusting automated systems has been found essential for adopting them [7, 48]. At the same time, trust research is somewhat controversial [15] because inappropriate trust may entail undesirable effects like misusing technology [11, 39]. In addition, trust is a complex topic. On the one hand, it has been defined in many different ways, depending on the field or context [36]. It also entails different themes like competence, benevolence, and reliance [7, 12, 13, 25, 34, 42]. On the other hand, it has been recognized that trust is not static but evolves [31, 44, 47]. Thus, measuring trust in automated systems is challenging and can be done either explicitly or implicitly.

Explicit measurements apply questionnaires or interviews to ask people about their trust perceptions. *One-dimensional* approaches measure trust with a single Likert-type question [31, 40, 44]. Although this method is quick and easy, it is susceptible to people interpreting 'trust' differently. Therefore, *multidimensional* approaches use Likert scales to measure trust as an ensemble of multiple constructs. For example, McKnight et al. [37] introduced the concept of *trusting beliefs* [57], consisting of the constructs *competence*, *benevolence*, and *integrity*. Later research added more constructs, including *perceived transparency* and *intention to return* [8, 49]. Overall, while a multidimensional approach is more nuanced than its one-dimensional counterpart, it requires longer questionnaires and is therefore more time-consuming.

Implicit measurements try to avoid the potential bias in self-reported explicit measurements by measuring trust through an intermediary. Examples are: loyalty measured by the number of logins after sign-up [38, 51]; accepting recommendations [13]; and time spent on a page, click-through rate, or page-exiting manner [20]. In the context of explaining recommender systems, implicit measurements for trust have not yet been widely adopted, possibly because intermediaries like loyalty require long(er)-term studies.

2.4 Trust in Explained Recommendations

Previous research has shown that providing explanations for recommendations can increase the acceptance of recommendations [13, 29], and increase people's trust in the recommender system [8, 48]. While previous studies typically focused on recommenders for movies or e-commerce, research in an e-learning context is limited [6]. This is unfortunate as Abdi et al. [2] recently demonstrated the potential of transparency for an educational recommender system: an Open Learner Model [10] improved understanding of and trust in recommendations for learning materials.

As trust is a relative measure, it must be compared to some baseline. Studies on the effects of explanations typically include a baseline with no explanations. However, a lesser applied baseline are *placebo explanations*: 'pseudo explanations' that are not semantically sensible [33], i.e., they do not reveal any information about why something was recommended, for example "*This has been recommended to you because this is what the algorithm calculated.*" Eiband et al. [17] found that placebo explanations can invoke similar trust levels as real explanations. Surprisingly, Nourani et al. [43] found conflicting results outside the domain of recommender systems: placebo explanations lowered the perceived accuracy of an image recognition system.

2.5 Research Gaps

Our literature overview shows that explaining recommender systems gains traction in the spirit of XAI and that studies often investigate how increased transparency affects trust in recommendations. However, we see two gaps. First, research on trust and explaining recommender systems often neglects adolescents: the presented literature primarily focuses on university students and young adults. Second, explaining recommender systems in an e-learning context is underexplored, which is unfortunate as e-learning platforms increasingly adopt recommendation algorithms [14, 28, 32, 56]. Our research addresses these gaps: we design hybrid explanations for an exercise recommender on an e-learning platform and investigate their effects on adolescents' *initial* trust (i.e., trust based on their first impressions of the platform).

3 MATERIALS AND METHODS

This section presents our e-learning platform with explanations for recommended exercises, our overall study design, and our data analysis decisions. Our research was approved by the ethical committee of *[Anon.]* (reference number *[Anon.]*).

3.1 E-learning Platform with an Exercise Recommender

For our study, we built upon an existing e-learning platform called Wiski [46], which was developed in Drupal, and contains over 1000 multiple choice exercises about mathematics topics in the *[Anon.]* high school curriculum. To tailor the difficulty level of exercises to students' level of mastery, we set up an *Elo rating system* [18] for students and exercises: if a student correctly solves an exercise, their Elo score rises and the exercise's Elo score drops, and vice versa. When picking exercises manually, students could thus estimate which exercises were relevant for them.

To also automate exercise selection, we implemented a recommender system in Python inspired by Dahl and Fykse [14]. Broadly, our recommender system combined Elo ratings and collaborative filtering: it looked for candidate exercises based on a student's Elo rating and recommended those that the student was most likely to answer correctly. More specifically, to recommend exercises for student A, our algorithm followed three steps. First, the 7 exercises closest to the value $\text{Elo}_\text{A} + 50$ were selected as candidates. We added the constant 50 to promote recommendations that slightly exceed the students' level of mastery [60]. Second, for each candidate exercise E, the algorithm estimated how many attempts A may need to solve E: it first looked for students who solved E, then used at most 40 nearest neighbors (Pearson similarity) to select those close to A in terms of attempts for previously solved exercises, and finally took a weighted average of their number of attempts for E. Third, the three candidate exercises with the lowest average number of attempts were recommended in ascending order.

(a) A real explanation for the REAL group

(b) A placebo explanation for the PLACEBO group

(c) No explanation for the NONE group

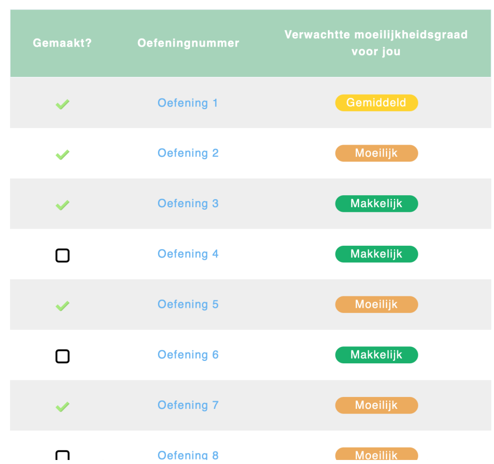

(d) Exercise list with difficulty labels (easy, average, hard)

Fig. 1. The three explanation interfaces in our randomized controlled experiment. In each interface, the top part (blue) shows real, placebo, or no explanations. The bottom part (green) allows users to return to the list of exercises from the same topic (d).

3.2 Explanations for Recommendations

Following a user-centered design process with think-aloud studies involving 12 participants (1 teacher, 5 middle school students, 6 high school students), we designed the three explanation interfaces in Figure 1 to accompany the recommended exercises. The first interface (Figure 1a) contained a real explanation, consisting of three parts [*English translation in brackets*]: (1) a why-statement which indicated that the exercise was recommended based on both the

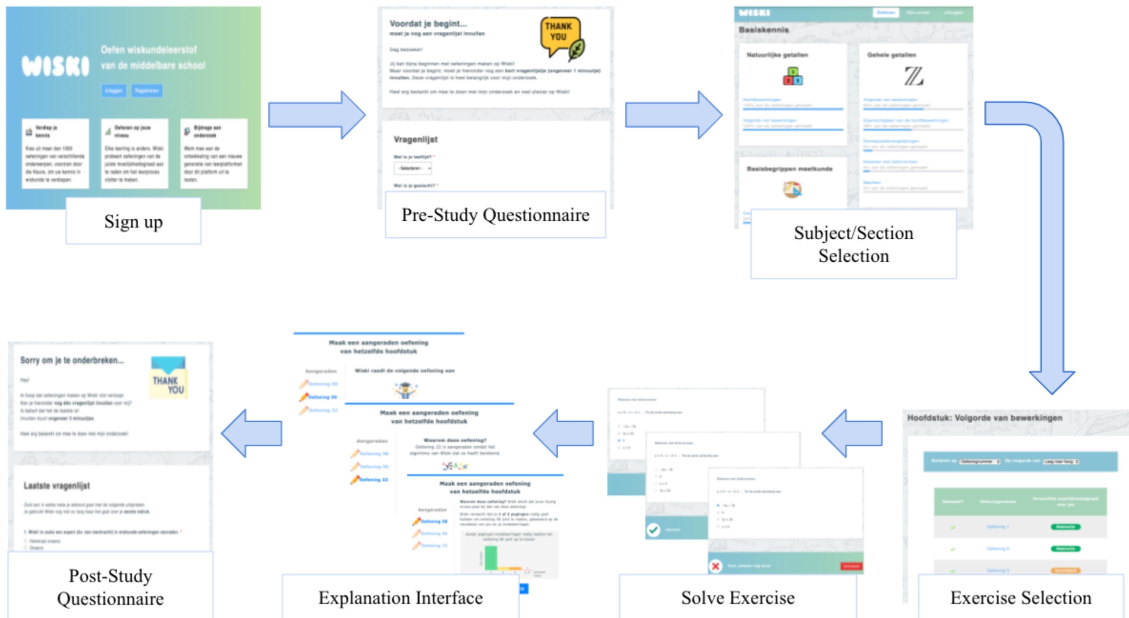

Fig. 2. Scheme of our study's flow: sign up, pre-study questionnaire, solving exercises and interacting with an explanation interface five times, and post-study questionnaire.

user's level of mastery and the exercise's difficulty [*Why this exercise? Wiski thinks that your current level matches that of this exercise!*]; (2) a justification-statement with the user's estimated number of tries needed to solve the exercise [*Wiski expects that you will need 1 or 2 attempts to answer exercise 4 correctly, based on your results and that of your peers*]; (3) a histogram of similar users' required tries for the exercise, inspired by Herlocker et al. [29] [*Number of attempts peers needed to solve exercise 4 correctly*]. The second interface (Figure 1b) contained a placebo explanation "*Exercise X is recommended because this is what Wiski's algorithm calculated*", which indeed conveys no information about how our recommendation algorithm works. Finally, the third interface (Figure 1c) simply stated that the exercise was recommended, without further clarification.

3.3 Participant Collection

We contacted teachers of 18 high schools in *[Anon.]* and invited them and their students to participate in our research. Teachers and students received an information leaflet that described the research process and stressed that students could not be coerced into participating and would receive an equivalent substitute task if they did not wish to participate. Interested students then gave informed consent, and students under the age of 16 also required signatures from their parents. In addition, we collected extra participants through snowball sampling.

3.4 Study Design

To assess the effect of our explanation interfaces on initial trust, we conducted a randomized controlled experiment [23] with three research groups: *REAL*, *PLACEBO*, and *NONE*, corresponding to the explanation interfaces in Figures 1a to 1c, respectively. Following the steps in Figure 2, all participants (1) registered on our platform and were randomly assigned

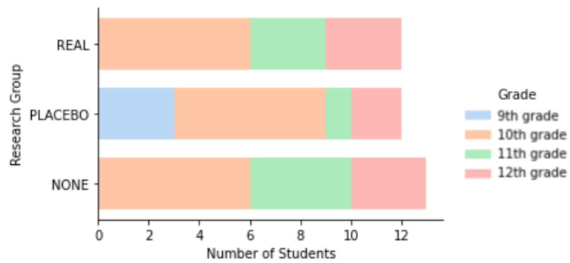

Fig. 3. Distribution of the 37 participating students over the three research groups.

a research group; (2) answered a pre-study questionnaire with questions related to their demographics, experience with computers and e-learning platforms, mathematical background, and self-perceived mastery in mathematics; (3) solved five exercises, and interacted with their research group's explanation interface after each exercise; (4) answered the post-study questionnaire in Table 2 with questions on trust; and (5) optionally used the platform freely until the end of the study. Thus, participants' experience on our platform only differed in the explanation interface shown after solving exercises. In the background, we also logged whether participants selected recommended exercises.

We decided to let participants answer the post-study questionnaire after five exercises because (a) they then all interacted with an explanation interface equally often, and (b) they often participated during a mathematics period at school and needed to finish in under an hour. The post-study questionnaire itself contained nineteen 7-point Likert-type questions divided into seven groups (see Table 2), and a text field after each group in which participants could motivate their responses. We measured trusting beliefs, consisting of *Competence* (Q1–Q5), *Benevolence* (Q6–Q8), and *Integrity* (Q9–Q11) with a validated questionnaire by Wang and Benbasat [7]. To fit the original questions in the scope of Wiski, we translated them to Dutch and made them easier to understand for adolescents by simplifying some vocabulary. The average of the scores for trusting beliefs, *Intention to return* (Q13–Q14), and *Perceived transparency* (Q15) yielded a *multidimensional trust* score. In contrast, *Trust* (Q12) assessed *one-dimensional trust* by explicitly asking about trust in Wiski's recommendations. Finally, *General questions* (Q16–Q19) collected extra information about how participants perceived explanations.

3.5 Data Analysis

We used non-parametric statistical techniques to avoid normality assumptions, similar to other work involving Likert-type data [e.g., 2, 13]. More specifically, we used the Mann-Whitney U test to check for significant differences between research groups, and Kendall's τ to test for correlations. To interpret the former as a difference in medians, we assumed equal data distributions in our three research groups. For all statistical analyses, we used Python 3.8.5 with libraries Pingouin 0.3.11 [54] and Seaborn 0.11.0 [59].

4 RESULTS

In total, 37 students (ages 13–18, 13 male, 24 female) participated in our research: 3 students were from 9th grade, 18 from 10th grade, 8 from 11th grade, and 9 from 12th grade. Figure 3 shows their distribution over the three research groups: 12 in REAL, 12 in PLACEBO, and 13 in NONE. Figures 4 and 5 show their responses to the post-study questionnaire.

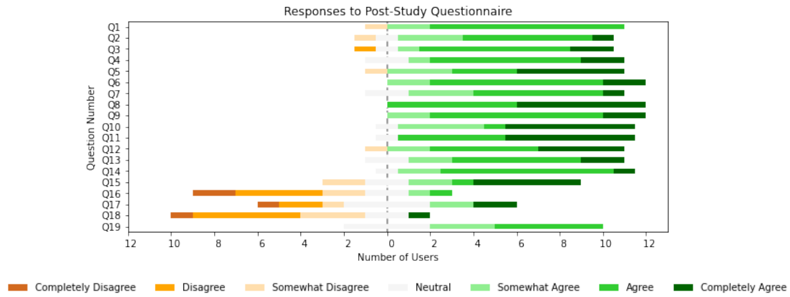

(a) REAL group

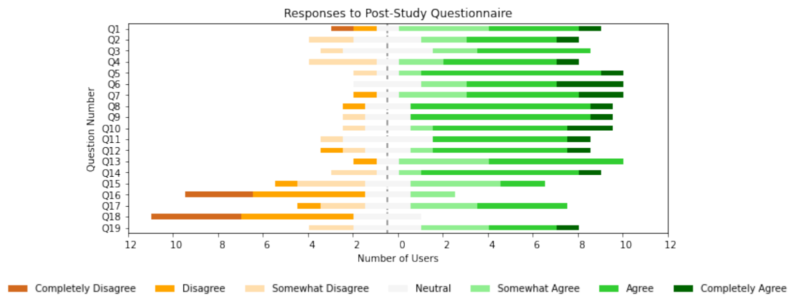

(b) PLACEBO group

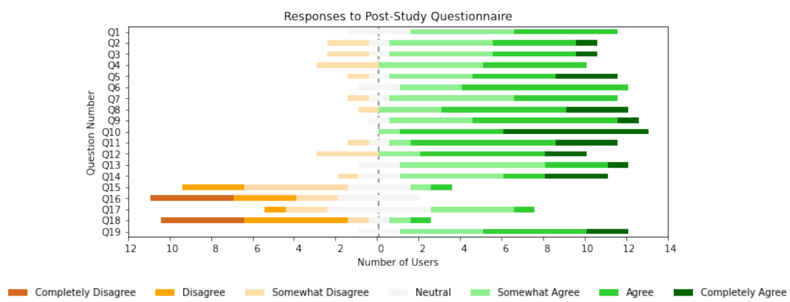

(c) NONE group

Fig. 4. Diverging bar charts of the responses to the post-study questionnaire for each research group.

4.1 Effects of Real Explanations

Tables 1a and 1b depict the outcomes of one-sided Mann-Whitney U tests, comparing REAL to NONE, and REAL to PLACEBO. Median competence, trusting beliefs, perceived transparency, and multidimensional trust were significantly higher in REAL ($p < 0.05$). However, there was no significant increase in one-dimensional trust or intention to return. For benevolence, there was only a significant increase ($p < 0.05$) when comparing REAL to NONE.

The qualitative responses on Q15 showed that perceived transparency was somewhat controversial in REAL. Some participants were positive about the explanations: "*I found the explanation that Wiski gave correct and satisfactory.*" Other participants did not seem to be satisfied with the explanations and may have wanted a different type of explanation:

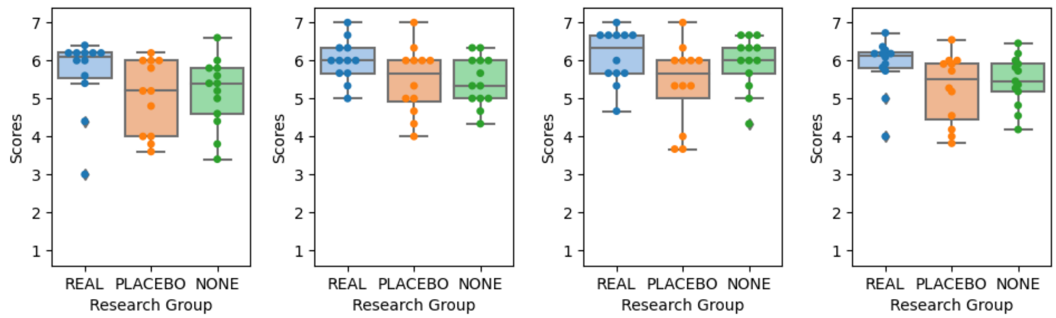

(a) Competence (b) Benevolence (c) Integrity (d) Trusting beliefs

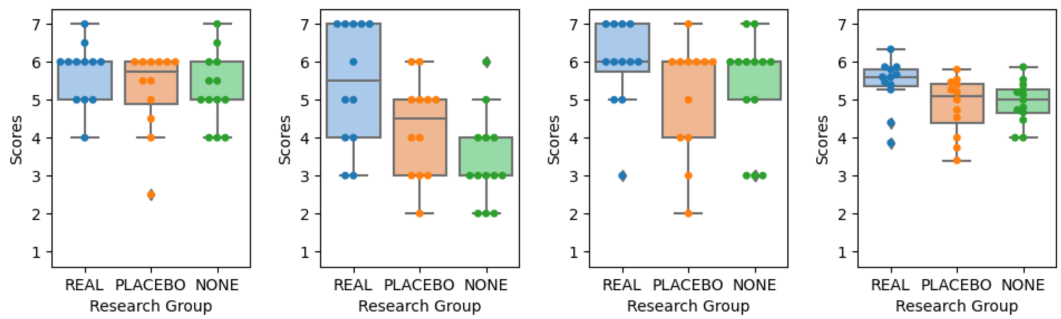

(e) Intention to return (f) Perceived transparency (g) One-dimensional trust (h) Multidimensional trust

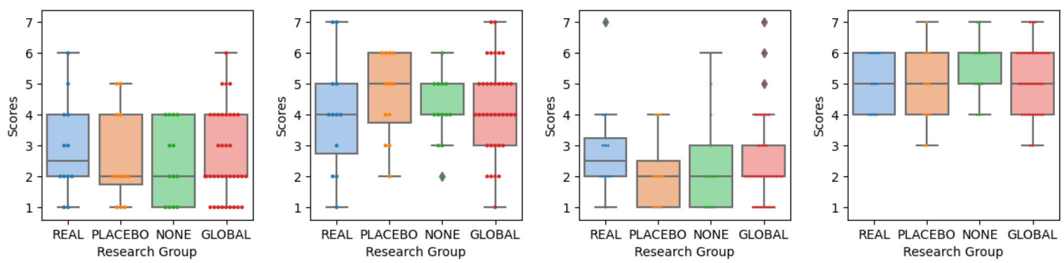

(i) Q16 (j) Q17 (k) Q18 (l) Q19

Fig. 5. Box plots of the responses to the post-study questionnaire.

"*Doesn't it just state how many tries Wiski thinks I would need to find the correct answer. It doesn't explain specifically.*" Finally, there was also evidence that some participants did not require explanations: "*I didn't really read the explanation...*"

4.2 Effects of Placebo Explanations

Two-sided Mann-Whitney U tests did not reveal any significant difference between the PLACEBO and NONE groups ($p < 0.05$): the two smallest p-values were 0.099 and 0.143 for perceived transparency and integrity, respectively; all other values were above 0.696. However, it is interesting that in our sample PLACEBO got the lowest median values for competence (Figure 5a) and integrity (Figure 5c).

Table 1. Results of one-sided Mann-Whitney U tests comparing the research groups. The common language effect size is the probability that a random value from one group is greater than a random value from the other group [61].

(a) REAL vs. NONE

	p-value	U value	CLES
Competence	0.029*	113.0	0.724
Benevolence	0.030*	112.5	0.721
Integrity	0.261	90.0	0.577
Trusting beliefs	0.038*	111.0	0.712
Intention to return	0.109	100.5	0.644
Perceived transparency	0.002**	130.5	0.837
One-dimensional trust	0.137	97.5	0.625
Multidimensional trust	0.014*	119.0	0.763

$^{*}p < 0.05$, $^{**}p < 0.01$, CLES = common language effect size

(b) REAL vs. PLACEBO

	p-value	U value	CLES
Competence	0.023*	37.5	0.740
Benevolence	0.074	47.0	0.674
Integrity	0.054	44.0	0.694
Trusting beliefs	0.030*	39.0	0.729
Intention to return	0.139	54.0	0.625
Perceived transparency	0.041*	42.0	0.708
One-dimensional trust	0.937	47.5	0.330
Multidimensional trust	0.013*	33.0	0.771

$^{*}p < 0.05$, CLES = common language effect size

As in REAL, the qualitative responses concerning perceived transparency (Q15) showed very different sentiments in PLACEBO. On the one hand, some participants did not perceive the placebo explanations as real explanations, as seen in responses like "*Wiski just says calculated by the algorithm of . . .*" and "*It would be nice for an extensive explanation as to why it is better to solve this exercise.*" On the other hand, other participants found the explanation satisfactory, as seen in responses like "*Wiski says that the algorithm recommends the next exercise thus I trust the algorithm*" and "*I don't think that there needs to be more explanation as to why an exercise has been recommended.*"

4.3 Effects of No Explanations

The qualitative responses on Q15 were quite consistent within NONE: close to all participants who gave a meaningful response indicated that they did not see an explanation or missed it. For example, one participant stated: "*I find it unfortunate that [the site] does not say why a certain exercise was recommended. It is nice to know why this exercise fits you, but there should also not be too much information as then it would not be fun to read.*" Yet, surprisingly, two participants seemed to believe they *did* receive explanations: "*If you want to solve a new exercise, it is useful that you know why this exercise is recommended, the website does this well*" and "*Yes I find that there is enough explanation.*" Finally, one participant formed a particular mental model of our recommender system: they thought recommendations depended on the self-reported mastery level of mathematics in the pre-study questionnaire.

4.4 Correlations

Figure 6 shows the correlations between the various trust constructs and one-dimensional trust: competence ($\tau = 0.68$) and integrity ($\tau = 0.72$) are correlated the most, whereas perceived transparency ($\tau = 0.16$) the least. In fact, perceived transparency has little to no correlation with any of the trust constructs. Figure 7 shows how all trust scores and questions Q16–Q19 are correlated. Especially notable is the moderate correlation between satisfaction with the level of recommended exercises (Q18) and most trust scores. We also found that one-dimensional trust is strongly correlated with trusting beliefs ($\tau = 0.673$) and multidimensional trust ($\tau = 0.624$).

4.5 Recommendation Clicks

Recall that the explanation interfaces recommended three exercises. Participants could either solve one of them next (i.e., accept a recommendation) or go back to the overview of exercises about the same topic and select an exercise

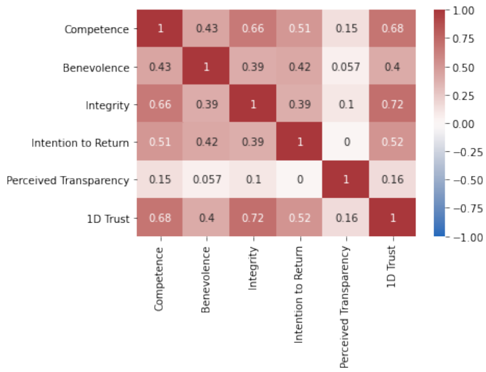

Fig. 6. Kendall's τ correlations between trust constructs and one-dimensional trust.

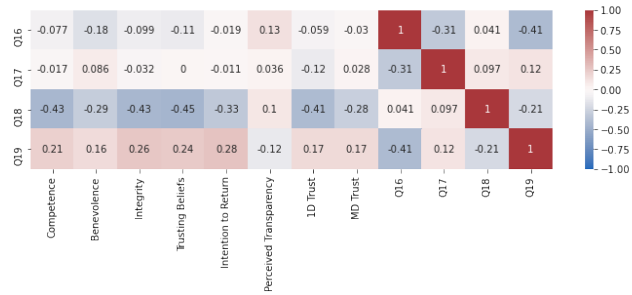

Fig. 7. Kendall's τ correlations between trust scores and questions on the need for explanations (Q16–Q19).

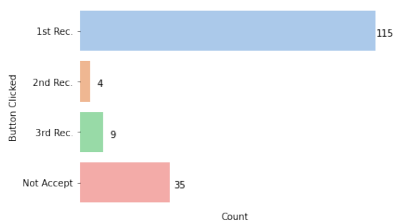

Fig. 8. Distribution of how often each option in the explanation interface was clicked.

themselves (i.e., not accept a recommendation, Figure 1d). Note that in our think-aloud studies, all participants were aware that our platform recommended three exercises, so we can assume that the same holds for our final study. Figure 8 shows that participants mostly decided to solve the first recommended exercise, followed by returning to the exercise overview. In addition, one-sided Mann-Whitney U tests revealed that the NONE group accepted less recommendations than both REAL ($p = 0.007$, $U = 67$, CLES = 0.827) and PLACEBO ($p = 0.039$, $U = 72$, CLES = 0.727).

5 DISCUSSION

This section answers our research questions by discussing how adding real, placebo, or no explanations to our e-learning platform affects adolescents' initial trust in our platform. Then, based on the observations, it underlines the need for tailoring explanations, and reflects upon the broader scope of explanations and recommendations in e-learning.

5.1 Explanations Increase Multidimensional Initial Trust...

Previous work has shown that well-designed explanation interfaces can increase adults' initial trust in a recommender system [17, 49, 62]. RQ1 asks whether the same holds for adolescents in an e-learning context. Two parts of our results suggest a confirmatory answer if trust is considered as an average of trusting beliefs, intention to return, and perceived transparency.

First, Table 1a shows that adding explanations significantly increased two out of three trust constructs: trusting beliefs and perceived transparency. The third construct, intention to return, was not significantly affected, which conflicts with the findings from Pu and Chen [49]: they reported that higher competence perception results in higher intention to return. One possible reason for this conflict might be that Pu and Chen's explanations assisted in buying expensive products, which seems more precarious than solving recommended exercises on an e-learning platform.

Second, participants with real explanations accepted more recommendations, that is, solved more recommended exercises. Building upon Cramer et al.'s [13] observation that acceptance of recommendations is correlated to trust, this further suggests that trust was higher for adolescents who saw real explanations.

5.2 ... But Not One-Dimensional Initial Trust

However, if trust is measured one-dimensionally with a single Likert-type question, there was *no* significant increase in trust compared to using placebo or no explanations. This shows that RQ1 cannot be answered in a univocal way, and puts our findings for increased trusting beliefs and multidimensional trust into perspective.

First, our results seem to imply that multidimensional trust measurements are more nuanced than their one-dimensional counterpart, which matches the well-known statement that trust is multi-faceted, and cannot be fully captured by a single question [30, 47].

Second, as most participants across the three research groups reported relatively high one-dimensional trust (see Figure 5g), the explanations may not have been the most important factor for trusting the e-learning platform. Instead, participants may have built initial trust mainly because of dynamically learned factors [30] like the perceived accuracy of the recommender system, the exercises' overall quality, or the platform's appearance. This is further backed by the correlations in Figures 6 and 7: whereas one-dimensional trust is barely correlated to perceived transparency and need for explanations (Q16, Q17, Q19), it *is* correlated to integrity, competence, and being satisfied with the exercises' level (Q18). Thus, explanations for recommendations seem to increase competence, which in turn increases initial trust. This further justifies the presence of competence in many definitions of trust [25, 42, 58].

5.3 Placebo Explanations Are a Useful Baseline

RQ2 is concerned with how placebo explanations influence adolescents' initial trust in our e-learning platform. We found no significant differences in initial trust when using placebo explanations over no explanations. This differs from Eiband et al.'s results [17], who found that placebo explanations *do* increase trust compared to no explanations. Reasons for the differing results could be the low sample size in both their and our study, the different study context, or the different methods for measuring trust. On a methodological level, Eiband et al. [17] suggest using placebo explanations as a placeholder when insufficient information is available for real explanations. Based on our results, however, we would discourage this as it may undermine the platform's perceived integrity (see Figure 5c).

However, when studying the impact of explanations, we do see several advantages for using placebo explanations as a baseline. For example, they allow to collect information about how critical participants stand towards explanations, and how attentive they are. In our study, we find it rather encouraging that most adolescents noticed that our placebo explanations were meaningless. Furthermore, placebo explanations allow to gain insights into how much transparency participants actually need. In our study, some adolescents required a more detailed explanation while others did not require much or any transparency. This underlines the importance of research on tailoring explanations based on transparency needs.

5.4 Tailoring Explanations Remains Important

Our qualitative data show that not all adolescents perceived the utility and transparency of our explanation interfaces in the same way. Some adolescents even had their own perception of what a good explanation is, and seek explanations that go beyond our focus on exercises' difficulty level and estimated number of attempts. To accommodate different transparency needs, it seems essential to tailor explanations to the audience that sees them.

On the one hand, the think-aloud studies during our user-centered design process gave us some insights into *what* parts of our real explanation interface could be tailored. First, middle school students (7th and 8th grade) typically found it harder to understand the histogram in our explanation, which suggests that this particular age group might require additional clarification for the histogram or an entirely different (visual) explanation. Second, some participants valued explicit wordings in the interface as it allowed them to process the given information quicker and better, while others considered this as rather redundant.

On the other hand, we can only speculate on *how* to concretize the tailoring process. One possibility is to give adolescents *direct* control over the explanations' type or detail level, or over whether they see any explanations at all. A potential drawback is that incomplete or no explanations can negatively impact adolescents' mental model of the recommender system, as illustrated by the participant in our NONE group who believed that the exercise recommendation depended on their self-reported mastery in mathematics. Another possibility to tailor explanations is to *indirectly* customize them according to personal characteristics [8, 40]. There is, however, an ethical challenge here as underage adolescents cannot or should not always pass delicate personality information without parental consent.

5.5 Taking a Step Back: Recommendations and Explanations in E-Learning

To conclude, we briefly reflect upon the premise of recommending exercises and explaining the underlying algorithm in e-learning. Do recommendations always need explanations? Should e-learning platforms always recommend exercises? We distinguish between situations in which little or much is at stake.

In *low-stakes* situations, accepting unsuitable recommendations does not have severe repercussions, so quickly accepting whichever recommendation seems reasonable. In our short-term experiment, students understood that accepting recommendations involved little risk, which may explain why they most often selected the first recommended exercise. In addition, some teachers instructed students to drill a specific topic, so it is plausible that some students were more interested in solving as many exercises as possible rather than carefully choosing their next exercise. In such 'drilling' situations, recommending only one exercise (the best fit) at a time might be sufficient, and full-fledged explanations might be excessive. However, in our experiment, students who were left in the dark as to why an exercise was recommended were more eager to select one themselves in the exercises overview. Perhaps this was the case because they perceived the displayed difficulty levels (see Figure 1d) as a kind of explanation. Thus, even in low-stakes contexts, it seems desirable to provide some minimal information about the (recommended) exercises.

In *high-stakes* situations, it becomes more important to investigate recommendations, and there, we hypothesize that explanations become more important too. When students have limited time to prepare for an exam, for example, it seems plausible that they seek a justification for why they should spend time solving a recommended exercise. Regarding the recommendation algorithm, we have three remarks: (1) in a school context, teachers are in the perfect position to judge which topics are best suited for a particular student, so it would be interesting to study how they can steer recommendations; (2) we believe it remains important to give students the freedom to select exercises themselves, for example to follow teachers' instructions; (3) contrary to our basic recommender system with one overall Elo score for each student, more sophisticated algorithms [1] could work with topic-specific Elo scores.

5.6 Limitations and Future Work

Our research has some limitations that affect the generalizability of our results. First, our sample size is relatively small: with only 37 participants divided over three research groups, our results should be interpreted cautiously. However, we do present a valuable data set of adolescents' initial trust in a recommender system, and our methods can be used as starting points for future research. Second, since Elo scores of students and exercises become more accurate as more students solve exercises, the accuracy of recommendations and explanations might have changed during the experiment. However, as participants were equally satisfied with the level of recommended exercises (Q18, see Figure 5k), this should not have biased the results significantly. Third, some participants communicated that the exercises on our platform are rather basic. If solving an exercise takes an insignificant amount of time, the importance of picking a suitable recommendation becomes smaller. Future studies could thus be conducted with more challenging exercises to investigate whether our results hold. Fifth, although the post-study questions for trusting beliefs were based on those by Wang and Benbasat [7], we modified and translated them to match them to an e-learning context and adolescents. Future work can validate our questionnaire. Sixth, our short-term study could only assess initial trust, whereas trust is an entity that evolves [31, 44]. In contrast, long-term studies could open up the possibility of measuring trust implicitly through loyalty [38, 51].

6 CONCLUSION

This paper tackled the complex topic of initial trust in an e-learning platform that explains why it recommends certain exercises. Specifically, we investigated how real and placebo explanations affect trust. Contrary to the vast majority of other human-computer interaction research on this topic, we focused on adolescents as the target audience.

Our randomized controlled experiment with 37 high school students showed that our explanation interface increases adolescents' initial trust when we measure trust as a multidimensional construct of trusting beliefs, intention to return, and perceived transparency. However, this effect did not hold when we used measurements of a single Likert-type question on trust. This two-sided result seems to imply that one question cannot capture the multi-faceted nature of trust, and that dynamically learned factors like perceived accuracy of the recommendation algorithm and the website's appearance may be the leading cause for gaining initial trust in the platform. Furthermore, compared to using no explanations, we found that placebo explanations did not offer any significant trust differences quantitatively. However, the divisive qualitative responses revealed that tailoring explanations based on transparency needs remains essential. Finally, we reflected upon whether explanations and recommendations are always desirable in e-learning and distinguished between low- and high-stakes situations.

In sum, while our study is rather small in sample size, our results do seem to indicate that using explanations on an e-learning platform is an asset for high school students. Therefore, accompanying recommendations with explanations should be considered when designing e-learning applications similar to ours for adolescents.

ACKNOWLEDGMENTS

We are very grateful to all students and teachers who participated in our studies.

REFERENCES

[1] Solmaz Abdi, Hassan Khosravi, Shazia Sadiq, and Dragan Gasevic. 2019. A multivariate elo-based learner model for adaptive educational systems. *arXiv* 1 (2019).

[2] Solmaz Abdi, Hassan Khosravi, Shazia Sadiq, and Dragan Gasevic. 2020. Complementing educational recommender systems with open learner models. *ACM International Conference Proceeding Series* (2020), 360–365. https://doi.org/10.1145/3375462.3375520

[3] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y Lim, and Mohan Kankanhalli. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–18.

[4] Amina Adadi and Mohammed Berrada. 2018. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). , 52138–52160 pages. https://doi.org/10.1109/ACCESS.2018.2870052

[5] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2020. Explainable Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58 (2020), 82–115. https://doi.org/10.1016/j.inffus.2019.12.012 arXiv:1910.10045

[6] Jordan Barria-Pineda. 2020. Exploring the Need for Transparency in Educational Recommender Systems. *UMAP 2020 - Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization* (2020), 376–379. https://doi.org/10.1145/3340631.3398676

[7] Izak Benbasat and Weiquan Wang. 2005. Trust In and Adoption of Online Recommendation Agents. *Journal of the Association for Information Systems* 6, 3 (2005), 72–101. https://doi.org/10.17705/1jais.00065

[8] Shlomo Berkovsky, Ronnie Taib, and Dan Conway. 2017. How to Recommend? User Trust Factors in Movie Recommender Systems. (2017), 287–300. https://doi.org/10.1145/3025171.3025209

[9] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: a visual interactive hybrid recommender system. In *Proceedings of the sixth ACM conference on Recommender systems*. 35–42.

[10] Susan Bull and Judy Kay. 2010. Open learner models. In *Advances in intelligent tutoring systems*. Springer, 301–322.

[11] Adrian Bussone, Simone Stumpf, and Dympna O'Sullivan. 2015. The role of explanations on trust and reliance in clinical decision support systems. *Proceedings - 2015 IEEE International Conference on Healthcare Informatics, ICHI 2015* October (2015), 160–169. https://doi.org/10.1109/ICHI.2015.26

[12] K. Chopra and W. A. Wallace. 2003. Trust in electronic environments. *Proceedings of the 36th Annual Hawaii International Conference on System Sciences, HICSS 2003* (2003). https://doi.org/10.1109/HICSS.2003.1174902

[13] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. *The effects of transparency on trust in and acceptance of a content-based art recommender*. Vol. 18. 455–496 pages. https://doi.org/10.1007/s11257-008-9051-3

[14] Ole Halvor Dahl and Olav Fykse. 2018. *Combining Elo Rating and Collaborative Filtering to improve Learner Ability Estimation in an e-learning Context*. Ph.D. Dissertation. Norwegian University of Science and Technology.

[15] Brittany Davis, Maria Glenski, William Sealy, and Dustin Arendt. 2020. Measure Utility, Gain Trust: Practical Advice for XAI Researchers. In *2020 IEEE Workshop on TRust and EXpertise in Visual Analytics (TREX)*. IEEE, 1–8.

[16] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).

[17] Malin Eiband, Daniel Buschek, Alexander Kremer, and Heinrich Hussmann. 2019. The impact of placebic explanations on trust in intelligent systems. *Conference on Human Factors in Computing Systems - Proceedings* (2019). https://doi.org/10.1145/3290607.3312787

[18] Arpad E Elo. 1978. *The rating of chessplayers, past and present.* Arco Pub.

[19] Daniel Fitton, Beth T. Bell, Linda Little, Matthew Horton, Janet C. Read, Michelle Rouse, and Nicola Toth. 2016. *Working with Teenagers in HCI Research: A Reflection on Techniques Used in the Taking on the Teenagers Project.* Springer International Publishing, Cham, 237–267. https://doi.org/10.1007/978-3-319-33450-9_10

[20] Steve Fox, Kuldeep Karnawat, Mark Mydland, Susan Dumais, and Thomas White. 2005. Evaluating implicit measures to improve Web search. *ACM Transactions on Information Systems* 23, 2 (2005), 147–168. https://doi.org/10.1145/1059981.1059982

[21] Fatih Gedikli, Dietmar Jannach, and Mouzhi Ge. 2014. How should i explain? A comparison of different explanation types for recommender systems. *International Journal of Human Computer Studies* 72, 4 (2014), 367–382. https://doi.org/10.1016/j.ijhcs.2013.12.007

[22] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.

[23] Rachel Glennerster and Kudzai Takavarasha. 2013. *Running randomized evaluations*. Princeton University Press.

[24] Michele Gorgoglione, Umberto Panniello, and Alexnder Tuzhilin. 2011. In CARS We Trust : How Context-Aware Recommendations Affect Customers ' Trust And Other Business Performance Measures Of Recommender Systems Michele Gorgoglione Umberto Panniello Alexander Tuzhilin. (2011).

[25] Tyrone Grandison and Morri Sloman. 2000. A survey of trust in internet applications. *IEEE Communications Surveys & Tutorials* 3, 4 (2000), 2–16. https://doi.org/10.1109/comst.2000.5340804

[26] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)* 51, 5 (2018), 1–42.

[27] David Gunning and David W. Aha. 2019. DARPA's explainable artificial intelligence program. *AI Magazine* 40, 2 (2019), 44–58. https://doi.org/10.1609/aimag.v40i2.2850

[28] Jonas Nilsberg Refsnæs Haldor Myre, Sondre Oldervoll. 2017. *Developing a web application for creating, solving, assessing and collecting data from interactive mathematical tasks in Python/Django, HTML5, Javascript, CSS and MySQL.* Ph.D. Dissertation. NTNU.

[29] J. L. Herlocker, J. A. Konstan, and J. Riedl. 2000. Explaining collaborative filtering recommendations. *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (2000), 241–250. https://doi.org/10.1145/358916.358995

[30] Kevin Anthony Hoff and Masooda Bashir. 2015. Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 57, 3 (May 2015), 407–434. https://doi.org/10.1177/0018720814547570

[31] Daniel Holliday, Stephanie Wilson, and Simone Stumpf. 2016. User trust in intelligent systems: A journey over time. *International Conference on Intelligent User Interfaces, Proceedings IUI* 07-10-Marc, 164 (2016), 164–168. https://doi.org/10.1145/2856767.2856811

[32] S. Klinkenberg, M. Straatemeier, and H. L.J. Van Der Maas. 2011. Computer adaptive practice of Maths ability using a new item response model for on the fly ability and difficulty estimation. *Computers and Education* 57, 2 (2011), 1813–1824. https://doi.org/10.1016/j.compedu.2011.02.003

[33] Ellen J. Langer, Arthur Blank, and Benzion Chanowitz. 1978. The mindlessness of ostensibly thoughtful action: The role of "placebic" information in interpersonal interaction. *Journal of Personality and Social Psychology* 36, 6 (1978), 635–642. https://doi.org/10.1037//0022-3514.36.6.635

[34] John D Lee and Katrina A See. 2004. Trust in Automation: Designing for Appropriate Reliance. *Human Factors* 46, 1 (2004), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

[35] Zachary C. Lipton. 2018. The mythos of model interpretability. *Commun. ACM* 61, 10 (2018), 35–43. https://doi.org/10.1145/3233231 arXiv:1606.03490

[36] Maria Madsen and Shirley Gregor. 2000. Measuring human-computer trust. In *11th australasian conference on information systems*, Vol. 53. 6–8.

[37] D. Harrison McKnight, Vivek Choudhury, and Charles Kacmar. 2002. Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research* 13, 3 (2002), 334–359. https://doi.org/10.1287/isre.13.3.334.81

[38] Sean M MCNEE, Shyong K LAM, Joseph A KONSTAN, and John RIEDL. 2003. Interfaces for eliciting new user preferences in recommender systems. In *Lecture notes in computer science*. Springer, Berlin, 178–187.

[39] Stephanie M. Merritt, Heather Heimbaugh, Jennifer Lachapell, and Deborah Lee. 2013. I trust it, but i don't know why: Effects of implicit attitudes toward automation on trust in an automated system. *Human Factors* 55, 3 (2013), 520–534. https://doi.org/10.1177/0018720812465081

[40] Martijn Millecamp, Katrien Verbert, Sidra Naveed, and Jürgen Ziegler. 2019. To explain or not to explain: The effects of personal characteristics when explaining feature-based recommendations in different domains. *CEUR Workshop Proceedings* 2450 (2019), 10–18.

[41] Sina Mohseni, Niloofar Zarei, and Eric D. Ragan. 2018. A Multidisciplinary Survey and Framework for Design and Evaluation of Explainable AI Systems. 1, 1 (2018). arXiv:1811.11839 http://arxiv.org/abs/1811.11839

[42] Bonnie M Muir. 1987. Trust between humans and machines. *International Journal of Man-Machine Studies* 27 (1987), 327–339.

[43] Mahsan Nourani, Samia Kabir, Sina Mohseni, and Eric D. Ragan. 2019. The Effects of Meaningful and Meaningless Explanations on Trust and Perceived System Accuracy in Intelligent Systems. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 7, 1 (2019), 97–105. https://ojs.aaai.org/index.php/HCOMP/article/view/5284

[44] Mahsan Nourani, Joanie T. King, and Eric D. Ragan. 2020. The Role of Domain Expertise in User Trust and the Impact of First Impressions with Intelligent Systems. (2020). arXiv:2008.09100 http://arxiv.org/abs/2008.09100

[45] Ingrid Nunes and Dietmar Jannach. 2017. A systematic review and taxonomy of explanations in decision support and recommender systems. *User Modeling and User-Adapted Interaction* 27, 3 (2017), 393–444.

[46] Jeroen Ooge. 2019. *Het personaliseren van motivationele strategieën en gamificationtechnieken mbv recommendersystemen.* Master's thesis. KU Leuven, Leuven, Belgium.

[47] Jeroen Ooge and Katrien Verbert. 2021. Trust in Prediction Models: a Mixed-Methods Pilot Study on the Impact of Domain Expertise. arXiv:2109.08183 [cs.HC]

[48] Pearl Pu and Li Chen. 2006. Trust building with explanation interfaces. *International Conference on Intelligent User Interfaces, Proceedings IUI* 2006, January 2006 (2006), 93–100. https://doi.org/10.1145/1111449.1111475

[49] Pearl Pu and Li Chen. 2007. Trust-inspiring explanation interfaces for recommender systems. *Knowledge-Based Systems* 20, 6 (2007), 542–556. https://doi.org/10.1016/j.knosys.2007.04.004

[50] Maxwell Szymanski, Martijn Millecamp, and Katrien Verbert. 2021. Visual, textual or hybrid: the effect of user expertise on different explanations. In *26th International Conference on Intelligent User Interfaces*. 109–119.

[51] Nava Tintarev and Judith Masthoff. 2011. Designing and Evaluating Explanations for Recommender Systems. In *Recommender Systems Handbook*. 479–510. https://doi.org/10.1007/978-0-387-85820-3

[52] Nava Tintarev and Judith Masthoff. 2012. Evaluating the effectiveness of explanations for recommender systems: Methodological issues and empirical studies on the impact of personalization. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 399–439. https://doi.org/10.1007/s11257-011-9117-5

[53] Chun-Hua Tsai and Peter Brusilovsky. 2019. Evaluating visual explanations for similarity-based recommendations: User perception and performance. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*. 22–30.

[54] Raphael Vallat. 2018. Pingouin: statistics in Python. *The Journal of Open Source Software* 3, 31 (Nov. 2018), 1026.

[55] Alfredo Vellido. 2020. The importance of interpretability and visualization in machine learning for applications in medicine and health care. *Neural computing and applications* 32, 24 (2020), 18069–18083.

[56] Katrien Verbert, Nikos Manouselis, Xavier Ochoa, Martin Wolpers, Hendrik Drachsler, Ivana Bosnic, and Erik Duval. 2012. Context-aware recommender systems for learning: a survey and future challenges. *IEEE transactions on learning technologies* 5, 4 (2012), 318–335.

[57] Giulio Vidotto, Davide Massidda, Stefano Noventa, and Marco Vicentini. 2012. Trusting beliefs: A functional measurement study. *Psicologica* 33, 3 (2012), 575–590.

[58] Y. Diana Wang. 2014. Building Trust in E-Learning. *Athens Journal of Education* 1, 1 (2014), 9–18. https://doi.org/10.30958/aje.1-1-1

[59] Michael L. Waskom. 2021. seaborn: statistical data visualization. *Journal of Open Source Software* 6, 60 (2021), 3021. https://doi.org/10.21105/joss.03021

[60] Kelly Wauters, Piet Desmet, and Wim den Noortgate. 2012. Item Difficulty Estimation: an Auspicious Collaboration Between Data and Judgement. *Computers & Education* 58 (2012), 1183–1193. https://doi.org/10.1016/j.compedu.2011.11.020

[61] K. L. Wuensch. 2015. CL: The Common Language Effect Size Statistic. http://core.ecu.edu/psyc/wuenschk/docs30/CL.pdf

[62] Yongfeng Zhang and Xu Chen. 2018. Explainable recommendation: A survey and new perspectives. *arXiv* (2018). https://doi.org/10.1561/9781680836592

A POST-STUDY QUESTIONNAIRE

Table 2. The questionnaire that participants answered at the end of the study. All questions were evaluated on a 7-point range. The group names in italics are for reference; participants did not see them.

No.	English original	Dutch translation
Competence		
Q1	Wiski is like an expert (for example, a teacher) for recommending math exercises.	Wiski is zoals een expert (bv. een leerkracht) in wiskunde-oefeningen aanraden.
Q2	Wiski has the expertise (knowledge) to estimate my math level.	Wiski heeft de expertise (kennis) om mijn wiskundeniveau te kunnen inschatten.
Q3	Wiski can estimate my math level.	Wiski kan mijn wiskundeniveau inschatten.
Q4	Wiski understands the difficulty level of math exercises well.	Wiski begrijpt de moeilijkheidsgraad van wiskunde-oefeningen goed.
Q5	Wiski takes my math level into account when recommending exercises.	Wiski houdt rekening met mijn wiskundeniveau om oefeningen aan te raden.
Benevolence		
Q6	Wiski prioritizes that I improve in math.	Wiski zet op de eerste plaats dat ik vorderingen maak in wiskunde.
Q7	Wiski recommends exercises so that I improve in math.	Wanneer Wiski oefeningen aanraadt, doet Wiski dat zodat ik vorderingen maak in wiskunde.
Q8	Wiski wants to estimate my math level well.	Wiski wilt mijn wiskundeniveau goed inschatten.
Integrity		
Q9	Wiski recommends exercises as correctly as possible.	Wiski raadt oefeningen op een zo correct mogelijke manier aan.
Q10	Wiski is honest.	Wiski is eerlijk.
Q11	Wiski makes integrous recommendations.	Wiski maakt oprechte aanbevelingen.
Trust (one-dimensional)		
Q12	I trust Wiski to recommend me math exercises.	Ik vertrouw Wiski om mij wiskunde-oefeningen aan te raden.
Intention to return		
Q13	If I want to solve math exercises again, I will choose Wiski.	Als ik nog eens online wiskunde-oefeningen maak, dan kies ik voor Wiski.
Q14	If I want to be recommended math exercises again, I will choose Wiski.	Als ik nog eens wiskunde-oefeningen aangeraden wil krijgen, dan kies ik voor Wiski.
Perceived transparency		
Q15	I find that Wiski gives enough explanation as to why an exercise has been recommended.	Ik vind dat Wiski genoeg uitleg geeft over waarom een oefening aangeraden is.
General questions		
Q16	I do NOT want any explanations about why an exercise has been recommended when I use Wiski.	Wanneer ik Wiski gebruik, wil ik GEEN uitleg over waarom een oefening wordt aangeraden.
Q17	I find receiving an explanation about why an exercise has been recommended more important than an explanation for why a movie has been recommended.	Ik vind uitleg krijgen over waarom een oefening wordt aangeraden belangrijker dan waarom een film wordt aangeraden.
Q18	I am NOT happy with the level of math exercises Wiski recommended.	Ik ben NIET blij met het niveau van de oefeningen die Wiski aanraadde.
Q19	I find it important to receive explanations when something (exercise/movie/product/...) has been recommended.	In het algemeen vind ik het belangrijk om uitleg te krijgen wanneer iets (oefening/film/product/...) wordt aangeraden.