

Designing an adaptive learner model for a mathematics game

MEES, Martyn <<http://orcid.org/0000-0002-9604-2277>>, JAY, Tim <<http://orcid.org/0000-0003-4759-9543>> and HABGOOD, Jacob <<http://orcid.org/0000-0003-4531-0507>>

Available from Sheffield Hallam University Research Archive (SHURA) at:

<https://shura.shu.ac.uk/22449/>

This document is the author deposited version. You are advised to consult the publisher's version if you wish to cite from it.

Published version

MEES, Martyn, JAY, Tim and HABGOOD, Jacob (2018). Designing an adaptive learner model for a mathematics game. In: CIUSSI, Melanie, (ed.) Proceedings of the 12th European conference on games based learning. Academic Conferences and Publishing International Limited, 800-807. [Book Section]

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Designing an Adaptive Learner Model for a Mathematics Game

Martyn Mees, Prof. Tim Jay, Dr. Jacob Habgood
Sheffield Hallam University, Sheffield, United Kingdom

M.Mees@shu.ac.uk

T.Jay@shu.ac.uk

J.Habgood@shu.ac.uk

Abstract

The RAIDING Project (Researching Adaptivity for Individual Differences in Number Games) aims to develop a game for 7-8 year olds, to develop their times tables and number bond skills. One of the design principles of the project is to implement a level of adaptivity into the game, so that the difficulty of the mathematical content adapts to the player's current level of arithmetic fluency. A learner model has been developed to enable the game to use previous gameplay performance to calculate the player's current level of arithmetic fluency, and thereby provide new tasks at an appropriate level of difficulty. A second design principle is to decouple the mathematical difficulty from gameplay rewards, so that progress in the game is achieved through time and effort rather than solely as a result of mathematical achievement. We predict that these two design principles will produce games that are motivating and help players to experience flow.

This paper describes and discusses our adaptive implementation, and our approach to decoupling of mathematical learning from rewards. We evaluate the success of the game to date and consider scope for potential development and improvement. We also show how the analytical data produced by the learner model has been used to identify unhelpful in-game behaviours and adapt the game design.

A future goal of the project is to explore whether the adaptivity of the learner model can be expanded to include gameplay ability (including elements hand-eye coordination and response times) and allow for separate dynamic adjustment of (non-mathematical) difficulty. We are particularly interested in investigating the affordances of such a "two-axis" flow in the game.

Keywords

game-based learning; adaptive difficulty; intrinsic motivation; arithmetic; flow

1. Introduction

Developing a child's basic numeracy skills is an important task in modern education. Without a solid foundation of skills, such as basic addition and multiplication, learning more complex mathematical concepts can be next to impossible for children. Using modern technology, such as video games and touchscreen devices, we can help a child develop and master these basic numeracy skills.

One problem we have identified in arithmetic number games is a lack of adaptivity to a player's level of arithmetical ability. Traditionally, most educational games link the difficulty of their learning content to progression in the game e.g. *Math Blaster* (Davidson & Associates et al., 1983-1999) or *Logical Journey of the Zoombinis* (TERC, 1996). The easiest learning content is delivered at the start of the game and progression to subsequent game content depends on demonstrating mastery of that learning content. Even previous mathematical games built around sound "intrinsically integrated" design principles link the learning content and game content in this way (Habgood and Ainsworth, 2011). This can potentially result in a game that does not closely match a player's level of arithmetic ability. The game can be too easy or too hard at various stages of gameplay, which means the player could end up either bored or frustrated. If a player is bored or frustrated, they will not wish to play the game in question, and the amount of learning that will occur will be lower than expected.

2. Related Work

Educational video games, including mathematical educational games, are not a new concept, with examples stretching back decades, such as *Oregon Trail* (MECC, 1985), *Math Blaster* (Davidson & Associates et al., 1983-

1999) and *Logical Journey of the Zoombinis* (TERC, 1996). How educational video games approach mathematics, however, is still a varying, fluid landscape. Some games subscribe to the classic "chocolate covered broccoli" approach (Bruckman, 1999), wherein the mathematics is confined to specific points throughout the game and has little connection to gameplay and strategy. In essence, these games could have been produced without the mathematics and have had a roughly identical gameplay experience. However, some mathematics games present a well-researched, pedagogically informed approach (Kebritchi and Hirumi, 2008). Well-designed titles can not only provide a fun experience, but also a marked increase in arithmetic ability.

One influential mathematics video game title, using a sound pedagogically informed game for research is *Zombie Division* (Habgood, J., 2005) The game utilises a design that focuses on intrinsically motivating players during the game, using a core mechanic which requires the player to attack numbered enemies with weapons that represent a certain divisor. Furthermore, the game is set within a motivating fantasy context: battling skeletons in ancient Greece. While the game design is entirely separate (and therefore could be implemented using any sort of theme, or potentially no theme at all), the fantasy context enhances the player's motivation while playing the game, and increased their time on task. Furthermore, competition between players to complete the most levels played an important role in the study. In terms of results that the studies showed, an intrinsically motivating game can improve arithmetic ability further than an extrinsically motivating game, as well as no game at all.

Adaptive difficulty is also not a new concept in commercial (non-educational) games. Various games, such as *Crash Bandicoot* (Naughty Dog, 1996), *Homeworld* (Relic Entertainment, 1999), *Flow* (Thatgamecompany, 2006) and many more have dynamic difficulty adjustment that causes the game to become easier or harder, depending on how well a player is performing in the game. In contrast, mathematics games have tended not to include dynamic difficulty, instead relying on static, ever increasing difficulty as the game progresses. However, more progress is being made in this field, and one such example of this is the e-learning environment *Math Garden* (Oefenweb.nl B.V., 2010). Developed by Oefenweb and the University of Amsterdam, the *Math Garden* software is a suite of maths mini-games in various fields, such as addition, subtraction, symbolic number representation, and more. As well as the maths component of the software, there is also a "metagame", where the player has to "care" for a garden, by playing the maths minigames, as well as earning "coins" during the maths games that can be used for purchasing decorative trophies in the game's menu system (Oefenweb.nl B.V., n.d.). The most interesting feature of *Math Garden*, however, is the adaptive difficulty system that is integrated into the software. Based upon a modified ELO ranking system (Klinkenberg, Straatmeier and van der Maas, 2011), the adaptive difficulty system takes into account not only how accurate a player is in terms of answering mathematical problems successfully, but by also taking into account how rapidly the player can answer those questions. After performing studies, their adaptive system was found to be reliable, accurate and valid in terms of "ranking" the players and providing them with mathematics problems of an appropriate difficulty. It also had good outcomes in terms of motivation, as children were playing the maths games not only in proper school hours, but also outside of those times, indicating that the game was enjoyable. Furthermore, initial level of arithmetic ability was not a factor in terms of enjoyment, with children with lower levels of arithmetic ability playing for roughly equivalent amounts of time as children with higher levels of ability (Klinkenberg, Straatmeier and van der Maas, 2011).

Math Garden was successful in improving the level of arithmetic ability for children over a period of time, with a significant improvement in ability across all the grades tested. The adaptive system provides a large amount of diagnostic material, in terms of the performance of the class in the various items available to play, as well as how often they played and when they played, both in terms of when in the day, as well as what days they played (Klinkenberg, Straatmeier and van der Maas, 2011).

3. The RAIDING Project

The RAIDING (Researching Adaptivity for Individual Differences in Number Games) project is developing a touchscreen game for Year 3 (7 to 8 year-old) primary school children. The game's design is based upon a novel approach to adaptive difficulty, which decouples the difficulty of the learning content from progress in the game (Mees, et al., 2017). During the early design meetings of the game, we determined that we required a learner model to allow for the adaptivity that we desired within the game. The design and development of the current learner model was developed based on a mature game design which had been subject to small scale,

informal trials. The trials were to not only ensure that our game design was satisfactory, but to also gather data on the distribution of response times for a range of correct answers, to ensure that the figures that we use inside the learner model are a good fit in terms of arithmetical ability of the target audience group.

3.1 The Game Design

The main goal of the game design was to make the mathematics both intrinsic to the gameplay, and a fun part of the game. To achieve this, we created two varieties of minigames. The first minigame is a times-table based game, where a large asteroid is mined to create "rocks" with a number attached to them. At the start of each minigame the player is told which times table to collect and if the number attached to a rock is part of it then the rock is sent to a "cart". If it is not part of it, it despawns with appropriately negative feedback. Once 30 rocks have been collected, the minigame is completed and the cart returns to a mothership. This provides a certain number of credits (based upon the amount of rocks collected), as well as a random number of "seeds" for planting at the central base. The second minigame is based around number-bonds and involves selecting two numbers that add up to a given value (such as $2+8 = 10$, $1+9 = 10$), using the same rock mining theme.

3.1.1 The Game Setting

One small aspect of the game's design that arose during our initial design meetings was the game's setting. We chose the space setting as it is a gender-neutral theme, with room to expand into various other secondary themes (as we did with the planting and alien systems). The space theme also translated well to a 2D, top down environment, with no complicated environments or physics needing to be created or programmed, as well as allowing for a simpler art style for the various assets that the game employs.

3.2 The Metagame

The metagame sits outside of the mathematical core-gameplay and allows players to customise their mothership utilising the credits and seeds that they have gathered during the course of gameplay. In particular, the seeds can be used to grow plants inside the mothership's biodomes. Growing plants not only acts as an outlet for creativity, but also provides the mechanism for space travel by converting fully grown plants into biofuels. Plants take anywhere from 2 to 8 minutes to grow in real time, and additional biodomes can be added to allow for more plants from different climates, as well as providing a home for aliens. Alien eggs are found on pre-set days during the first 7 days (the player gets a single alien on days 3, 5 and 7) and can then randomly acquire aliens at any point after that. These eggs, require a day (in real time) to hatch. This acts as a "return trigger" to the game. Return triggers are events that make the player want to return to the game. Games can employ multiple types of return trigger, such as the appointment triggers (return at a set time for a reward), competitive triggers (such as leader boards, player vs player content, etc.), and social commitment triggers (when a player's response is needed for another player to continue) (Luton, 2013). Our seed growth and alien hatching mechanics act as appointment triggers for the game. Furthermore, once an alien is hatched, they require a certain variety of plants in their biodome to remain "happy" motivating the player to find more seeds and grow more plants.

The metagame requires credits and seeds, which are obtained from the core loop of repeating maths challenges. Progression in the game therefore requires mathematical minigames to be completed, but the progression of the mathematical content is controlled separately.

3.3 The Learner Model

One of the primary reasons for developing a learner model was that the game needed a way of recording the player's mathematical performance and arithmetic ability for each minigame as well as over the entire game-playing session. Furthermore, the game required a way of using this data to provide a set of probabilities for the game's random number generator to determine the type of game, the level of arithmetic difficulty, and which numbers to select. This should focus on the weaker parts of player's number knowledge which allow the player to further develop their arithmetic ability, without being too difficult and discouraging. Therefore, we developed the game's learner model.

The learner model is a vital part of the game, as it determines both the type of mathematical problems to generate for the player, and the precise probability of each number occurring. To do this, it has to calculate the

association strength of each individual "number fact" (such as 2 times 3, or 3 plus 7 equals 10) and the overall performance on each times-table or set of number bonds that the number fact belongs to.

Let us consider an example of how the system works for a times table minigame. Assuming this is the first time this game has been played, the learner model will provide a simple, even distribution of probability for each number in the times table. Rocks are generated in the minigame, and when a player taps on one of these rocks, the game records i) what rock was tapped, ii) whether or not it was a correct or an incorrect selection, iii) what number the rock had attached to it, and iv) the time the rock was spawned and selected in game. The same data is also recorded for rocks that haven't been tapped, and then have despawned after their lifespan has run out (which takes 10 seconds for a times-table game). This data is not only stored internally for the learner model, but also as an external log for the analytics system so that we can analyse the raw game data afterwards right down to a rock by rock level.

The minigame ends either when the player completes it by selecting thirty correct rocks, or by manually terminating the game session via the provided exit button. At this point the game processes the data that has been collected, exporting the captured data to the filesystem, as well as providing the learner model with the vital aggregate statistics for the recently completed game.

The learner model uses the data for every rock to calculate how many correct and incorrect answers there were overall for the minigame, as well as on an individual number fact level. Table 1 describes what is regarded as a correct and incorrect answer, for both individual numbers, as well as the entire times table in that game session. Furthermore, the average time to answer correctly is calculated for each individual number fact. Once this is done, the data is added to the learner model and the association strength is calculated.

Table 1: Required Responses for Correct and Incorrect Answers in the Times Table Minigame

	Correct Answer	Incorrect Answer
Overall Times Table	Rock in the times table was selected. Rock that wasn't in the times table despawned.	Rock that wasn't in the times table was selected. Rock that was in the times table despawned.
Individual Number Fact	Rock that was this number was selected.	Rock that was this number despawned.

Adding this new data to the learner model, and calculating the association strength involves multiple stages. Adding the data to the learner model, involves not only adding the data to a number of internal lists, but also calculating a set of weighted rolling averages, based upon this new data.

The weighted rolling averages for the player's accuracy and the player's average correct answer time are calculated first. The rolling averages take into account the player's last five games (including the game that has just been played), and the newest data that has been added to the learner model is weighted more heavily, than the older data. Figure 1 shows the formula used for calculating the weighting.

To calculate the weighting for a data value, this formula is used:

$$\frac{\text{Entry Position in List} + 1}{\text{Number of items in List} \times (\text{Number of items in List} + 1) \div 2}$$

The entry position is the internal index, which starts at zero. The newest game is the last item in the list.

So, for example, if there were 5 games of data:

- Oldest: $0+1 / (5 \times (5 + 1) / 2) = 1/15$
- $1+1 / (5 \times (5 + 1) / 2) = 2/15$
- $2+1 / (5 \times (5 + 1) / 2) = 3/15$
- $3+1 / (5 \times (5 + 1) / 2) = 4/15$
- $4+1 / (5 \times (5 + 1) / 2) = 5/15$

Figure 1: Rolling Average Formula

We chose to use a weighted rolling average system like this one because we felt that more recent games are more relevant in terms of analysing the player's performance, whereas the data for previous games becomes less relevant over time. This averaging process also helps to smooth out any atypical performance, while still taking into account that game.

Once these rolling averages have been calculated, the next stage is to calculate the association strength for each number. The formulae used to calculate the association strength are detailed in figure 2.

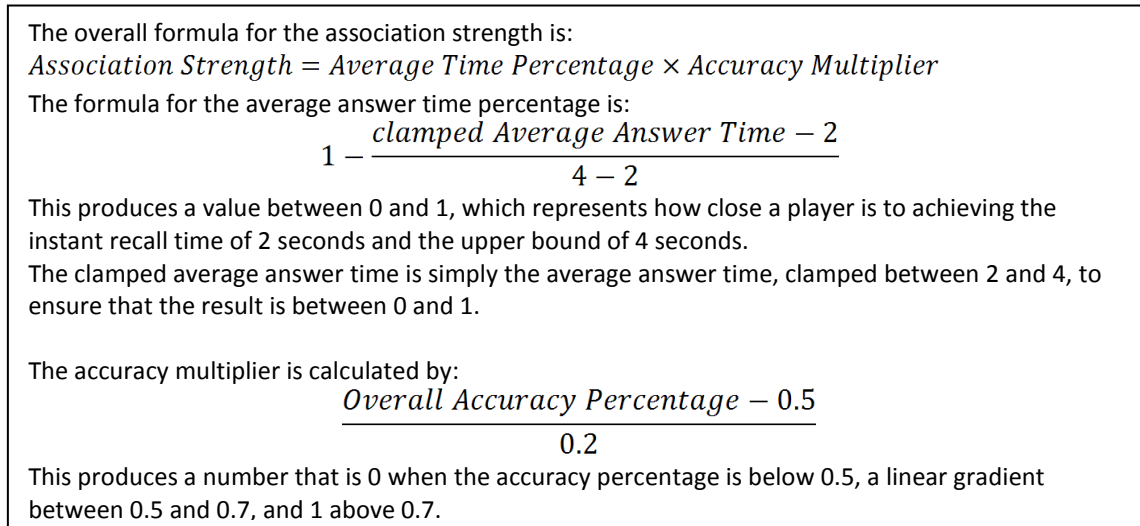


Figure 2: Association Strength Formulae

The average time to answer correctly is used as a measure for association strength, in line with Ashcraft (1982). We make the assumption that an arithmetical question answered correctly in 2 to 2 ½ seconds is indicative of a recall strategy being used. If the player is taking more than 4 seconds to answer, then the player is more likely to be using a manual method to find the answer, such as counting on their fingers, or calculating it in some way.

The next part of determining the association strength involves calculating the accuracy multiplier (figure 2). The formula means that if a player's accuracy is lower than 50%, then the association strength is regarded as zero. If the accuracy is greater than 70%, then we assume that the player knows that number fact, and then the only thing we take into account is how quickly they answered it. In between these two percentages, we calculate a multiplier to linearly adjust the association strength as their accuracy increases.

Once the association strength has been calculated for each individual number fact, the association strength is calculated for the entire times table. The same method applies for the overall times table, except that the average time to answer correctly is calculated by averaging the average times from each individual number fact.

The final step for the learner model is to calculate the thresholds for unlocking the next set of times tables. The learner model has three tiers of times table: tier 1 is the 10, 2 and 5 times tables; tier 2 is the 3 and 4 times tables; and tier 3 is the remaining 6,7, 8 and 9 times tables. For the game to advance to the next tier, all times-tables in the tier below must have an association strength of at least 0.75, and the player must have selected at least 30 rocks in each times-table in that tier. These values were chosen as they allow players with a good level of accuracy and average correct answer time to proceed in the game before they become tired of the same content.

The process that occurs when the player is playing a number bonds game is mostly the same, with a few exceptions. One exception is that the number facts for number bonds involve a pair of numbers, such as 1 and 9, or 3 and 7. We regard these pairs as commutative, as in game the pair can be selected in any order, and therefore it doesn't matter how they are stored. The next difference arises when calculating correct and incorrect answers, due to the fact that number bonds require a pair. This difference has been noted in table 2.

Table 2: Required Responses for Correct and Incorrect Answers in the Number Bonds Minigame

	Correct Answer	Incorrect Answer
Overall Times Table	A rock was correctly selected as part of a pair. A rock was correctly ignored as it couldn't make a pair.	A rock was ignored, but at some point could have made a pair. A rock was selected as part of a pair, but it was incorrect.
Individual Number Fact	A rock was correctly selected as part of a pair.	A rock was selected as part of a pair, but it was incorrect.

The final difference is in the different tiers that number bonds use. Tier 1 is simply number bonds to 10, tier 2 is number bonds to 20, and tier 3 is number bonds to 15.

3.4 Data Driven Design

The game was subject to a number of formative trials with children during its development, which were used to inform certain elements of the game design. Both the mothership construction and alien features were added after these early trials suggested the meta-game required more depth. Children were always requesting aliens in their feedback, so these became a central reward behind the metagame

However, to fully explore the effectiveness of the game and its learner model we are currently running a number of studies in local primary schools. We are particularly interested in exploring whether the learner model works, and whether the game is effective, and can hold the attention of the children over a sustained period. The final results of these studies are not yet available, but some preliminary data is explored below relating to the learner model.

We ran a study in a primary school in the Sheffield area, with the school's two classes of year 3 (7-8 years old) pupils. The classes were of mixed ability and had roughly the same number of pupils. We had one class play the game over a two-week period, for 20 minutes a day, with the other acting as a control, and then swapped which group was the control and which played the game for the next two weeks.

During this study, the game saved the state of the learner model at the end of each day, to allow us to examine the internal data of the learner model, as well as perform additional analysis and investigation (Habgood, Jay and Mees, 2018). Tools were also written to take this data, in the same format as the saving/loading system in the learner model, and output .csv spreadsheets with the data in a format that allows for easy graphing of a player's accuracy and average correct answer time. An example of these graphs is provided as Figure 3.




Figure 3: Graphs from Learner Model Data

Utilising these graphs, alongside observations from the study, we uncovered a curious behaviour that occurred towards the end of the study. We noticed that some of the charts (an example of which is the charts that make up Figure 3) showed that the accuracy was dropping dramatically towards day 3 in week 2, while the average

time to answer correctly was rapidly decreasing to below two seconds. Furthermore, we had noticed that towards the end of the study, a small minority of the children had developed a tactic where, they began randomly tapping on every rock that appeared in the game, to more rapidly advance in the metagame. As it currently stands, the game has no real penalty for this tactic beyond visual/aural effects, and this must be investigated further.

4. Future Improvements and Developments

The learner model appears to be effective as a tool for calculating a player's current level of arithmetic ability and generating appropriate number problems to allow the player to get better at their arithmetic. However, there are improvements which could be made, and these are now discussed alongside potential future developments for the learner model.

One of the issues with the learner model is how the "tiers" system limits the granularity of advancement. A player can be on the cusp of achieving a high enough level of association strength for it to unlock the next tier of times table, but just might not have fast enough reaction times to allow the next set to unlock. This can end up frustrating a player, as they may know the answers, but end up being stuck on a times table and unable to progress because they physically cannot react quickly enough for the game to proceed.

Based upon some of our results, one hypothesis for the effectiveness of the game is that the game engenders a feeling of flow on the player. Flow, in its most basic sense, is the feeling of the loss of the tracking of time and self-consciousness (Chen, 2006). Flow is commonly represented on two axes, where challenge and ability must be balanced in order to achieve a flow state (Csikszentmihalyi, 1975).

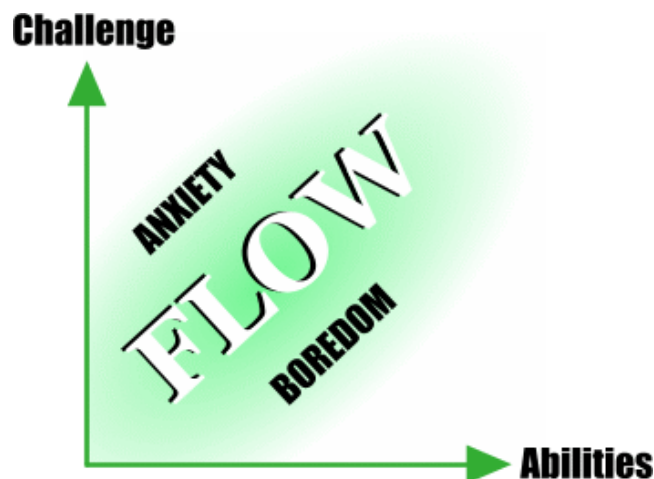


Figure 4: A classic representation of Csikszentmihalyi's "flow channel".

This element is an important one in video games, and indeed, forms the major hypothesis of how our game develops arithmetic ability over a period of time. With the decoupling of the arithmetic difficulty from the game difficulty (in terms of object moving speed, time that the rocks remain on screen, spawn rate, etc.), this allows for the player to achieve equal progress in the meta-game, regardless of the player's current level of mathematical ability.

One further development that we are exploring is the idea of adding a second "dimension" of flow to the game, based upon the gameplay difficulty, and adapting that to an individual player's type of gameplay and learning, to allow for the maximum level of learning and flow for a player. This would involve numerous modifications to the existing analytics and learner model paradigm, as we would have to track more variables, especially in terms of gameplay, and then also calculate which of these variables are more important to a particular player, which styles of gameplay does a player not like, and then have the game respond to this by adjusting various mini-game and metagame elements, and continue this feedback until the game is in a "comfortable" place for the player. During this, variants of the games would still be included, to allow for the player to have a different style of game, to continue to adapt the player's game style by gathering more data, and to continue to help maintain the sense of flow, and not let the player experience boredom.

5. Acknowledgements

This work was made possible with the help of a Research Project Grant from the Leverhulme Trust. We are grateful to the Leverhulme Trust for their support and to the children and school who have participated in the research.

References

- Ashcraft, M. H., (1982), "The development of mental arithmetic: A chronometric approach". *Developmental review*, 2(3), 213-236.
- Bruckman, A., (1999), "Can educational be fun?", *Game developers conference*, issue 99, pp. 75-79.
- Davidson & Associates/Knowledge Adventure, (1983-1999), *Math Blaster*, [Video Game] Apple II, DOS, Windows, Macintosh, Torrance, California: Davidson & Associates/Knowledge Adventure
- Habgood, J., (2005), *Zombie Division*, [Video Game] Windows, University of Nottingham
- Habgood, M.P.J., (2007), *The Effective Integration of Digital Games and Learning Content*, PhD thesis, University of Nottingham.
- Habgood, M.J., Ainsworth, S.E., 2011, "Motivating children to learn effectively: exploring the value of intrinsic integration in educational games", *Journal of the Learning Sciences*, vol. 20, issue 2, pp. 169-206.
- Habgood, J., Jay, T., Mees, M., (2018), *Unpacking Effective Learning Through Game Analytics*, In: CHI 2018 Data-driven educational game design, Carnegie Mellon University.
- Kebritchi, M., Hirumi, A., (2008), "Examining the pedagogical foundations of modern educational computer games", *Computers & Education*, vol. 51, issue 4, pp. 1729-1743
- Klinkenberg, S., Straatmeier, M., van der Maas, H.L.J., (2011), "Computer adaptive practice of Maths ability using a new item response model for on the fly ability and difficulty estimation", *Computers & Education*, vol. 57, issue 2, pp. 1813-1824.
- Lee, J., Luchini, K., Michael, B., Norris, C., Soloway, E., (2004) "More than just fun and games: assessing the value of educational video games in the classroom", *CHI '04 Extended Abstracts on Human Factors in Computing Systems*, pp. 1375-1378.
- MECC, (1985), *The Oregon Trail*, [Video Game] Apple II, Brooklyn Center, Minnesota: MECC
- Mees, M., Jay, T., Habgood, J., Howard-Jones, P., (2017), "Researching Adaptivity for Individual Differences in Number Games", *CHI PLAY '17 Extended Abstracts: Extended Abstracts Publication of the Annual Symposium on Computer-Human Interaction in Play*, pp. 247-253.
- Naughty Dog, (1996), *Crash Bandicoot*, [Video Game] PlayStation, Tokyo, Japan: Sony Computer Entertainment
- Oefenweb.nl B.V., (2010), *Math Garden*, [Video Game] Web browser, Amsterdam, Netherlands
- Oefenweb.nl B.V., (n.d.), *Math Garden More Info*, viewed 1st May 2018, <https://www.mathsgarden.com/more-info/>
- Relic Entertainment, (1999), *Homeworld*, [Video Game] Windows, Bellevue, Washington: Sierra Interactive
- TERC, (1996), *Logical Adventure of the Zoombinis*, [Video Game] Windows, Macintosh, Eugene, Oregon: Brøderbund Software, Inc.
- Thatgamecompany, (2006), *Flow*, [Video Game] PlayStation 3, Tokyo, Japan: Sony Computer Entertainment