



AI in Education, Learner Control, and Human-AI Collaboration

Peter Brusilovsky¹

Accepted: 7 June 2023 / Published online: 21 August 2023
© International Artificial Intelligence in Education Society 2023

Abstract

User control and human-AI collaboration are two related directions of research in the modern stream of work on human-centered AI. The field of AI in education was an early pioneer in this area of research, but now it lags behind the work on user control and human-AI collaboration in other areas of AI. This paper attempts to motivate further research on learner control and human-AI collaboration in educational applications of AI by presenting a review of the current work and comparing it with similar work in the field of recommender systems.

Keywords Artificial intelligence · Human-centered AI · User control · Human-AI collaboration

Introduction

In recent years, the use of Artificial Intelligence (AI) technologies has expanded to many areas where they directly affect the lives of many people. AI-based approaches advise human decision-makers whether it is a good time to discharge a patient from a hospital, who should be released on bail, and whether a specific student is at risk of failing a course. The increased reliance on AI came with a variety of problems and motivated a rapid rise in research on “human-centered AI” (Shneiderman, 2022), which attempted to address and minimize the negative effects of using AI technologies. Among the ideas of human-centered AI is *user control* - engaging users in AI decision-making to improve the results and prevent possible errors and biases.

The field of AI in Education was among the first to explore the ideas of user control. Most importantly, early work on open and editable learner models in computer assisted language learning and intelligent coaching systems (Bull, 1993; Bull et al., 1995; Cook & Kay, 1994; Kay, 1997) explored the opportunity for learners to collaborate with AI in the learner modeling process, that is, to visualize and maintain

✉ Peter Brusilovsky
peterb@pitt.edu

¹ School of Computing and Information, University of Pittsburgh, Pittsburgh, PA, USA

the usually hidden content of learner models used by most personalized learning systems. Research on open learning models along with pioneering research on cooperative user models and open user profiles (Kay, 1994; Waern, 2004) were critical to shaping the modern stream of research on open user models in other types of personalized systems (Ahn et al., 2015; Bakalov et al., 2013; Glowacka et al., 2013). However, the modern stream of work on human-AI collaboration and user control in a broader field of AI produced almost no follow-up in research on AI in Education (AIED). Does it mean that user control and human-AI collaboration have no value in applications of AI in education beyond classic open learner modeling?

This paper attempts to answer the above question by demonstrating a range of opportunities for human-AI collaboration and user control in AIED and illustrating these opportunities through a set of examples. However, before proceeding to the main content of the paper, the author wants to make two important comments on some terminology used in the paper. First, it could be useful to distinguish user control and human-AI collaboration as two different ways to implement human-centered AI that could be traced back to early research in the field. The term “user control” stresses that the user is the senior partner, the one in charge. While AI could do main “weight lifting”, for example, processing a large body of information, humans have tools to examine, control and adjust the way AI operates (Kay, 1994). In contrast, the term “human-AI collaboration” stresses that humans and AI are equal partners in achieving the overall goal, and each party contributes according to its strengths and weaknesses (Bull, 1993).

Second, in the context of this paper focused on AI in education, the term “user control” might be confusing, since intelligent educational systems have at least two distinct categories of users: teachers and learners. Since this paper is focused on learners as users of these systems, in the paper I will use a more specific term *learner control*. Note that the general ideas of learner control are extensively explored in the field of education, in particular, forming one of the foundations of increasingly popular self-regulated learning (Bjork et al., 2013). As a result, the term “learner control” is already widely used in the field. Surprisingly, learner control is frequently positioned as an alternative to AI-based learning approaches such as personalized learning, creating the questionable assumption that AIED and learner control are not compatible with each other. This paper, however, talks about learner control in the field of intelligent educational systems, a prospect extensively discussed in (Kay, 2001). Within this field, research on learner control focuses not on a broader idea of learners’ control over their learning, but on learner control over AI technologies integrated into the learning process. As this paper attempts to demonstrate, learner control over AI technologies is not only possible, but could also increase the value of these technologies.

Personalized Content Selection

To make a case for human-AI collaboration and learner control in AIED, I focus on a group of AIED technologies frequently referred to as *personalized guidance* or *personalized content selection*. This group of technologies focuses on guiding learners to the most appropriate learning content by considering their current level of knowledge, interests, and other factors. Specific technologies within this

group – adaptive sequencing (McArthur et al., 1988), personalized course generation (Diessel et al., 1994), adaptive navigation support (Brusilovsky, 2007), adaptive presentation (Bunt et al., 2007), and recommendation of learning content (Drachslers et al., 2015) – explore different ways to achieve this general goal. Focusing on one group out of many AI technologies used in the learning process helps to present a more systematic case. Moreover, focusing specifically on personalized content selection enables me to connect the opportunities for learner control and human-AI collaboration in AIED with opportunities for user control in a related area of recommender systems where research on user control is rapidly expanding. In turn, it helps to build bridges between learner control in AIED and user control in “big AI”.

To demonstrate the opportunities for learner control in personalized content selection, it is helpful to understand the main components of this process. The process usually starts with the learner model, which represents user features that are essential for content selection. Frequently, it means the user’s current knowledge level, but in some cases, it could also be needs, interests, preferences, etc. (Brusilovsky & Millan, 2007). Based on the current state of the model and the current context, the content retrieval engine selects the most relevant content. In the final stage, this content is presented to the user in some way. The nature of the retrieval engine and the form of content presentation could be different for different content selection technologies. For example, a content-based recommendation engine selects relevant content based on concepts and skills that this content enables learners to practice, while a social navigation support mechanism focuses on the performance of learners who worked with this content before. As past research demonstrates, learner control and human-AI collaboration could be applied to each of these three steps. In the following three sections, I provide examples of doing it at each of these steps.

Learner Modeling

Learner modeling is the starting point for all kinds of adaptation and personalization in AIED, including personalized content selection. Not surprisingly, engaging the learner in understanding, building and controlling the content of the learner model through an open learner model (OLM) is the most popular approach to control or collaborate with AI in AIED. It was an important early discovery in the field of AIED and user modeling (Bull, 1993; Bull et al., 1995; Cook & Kay, 1994) and, as mentioned in the introduction, it was one of the first examples of user control across all kinds of intelligent systems. A direct analog for it in the “big AI” are various open user profiles and user models, which are used in personalized search and recommender systems (Ahn et al., 2015; Bakalov et al., 2013; Glowacka et al., 2013; Kay, 1994; Waern, 2004). In this case, it could be argued that it was the success of OLM in AIED that motivated research on open user modeling in other types of AI applications. Today, open learning models (OLMs) are considered one of the core AIED approaches. OLMs were promoted by several highly cited review papers and used in a broad range of AIED systems (Bull & Kay, 2007; Dimitrova et al., 2007; Bull, 2020). The use of OLM in the form of *skillometers* in cognitive tutors (Corbett et al., 2000) made OLM a popular

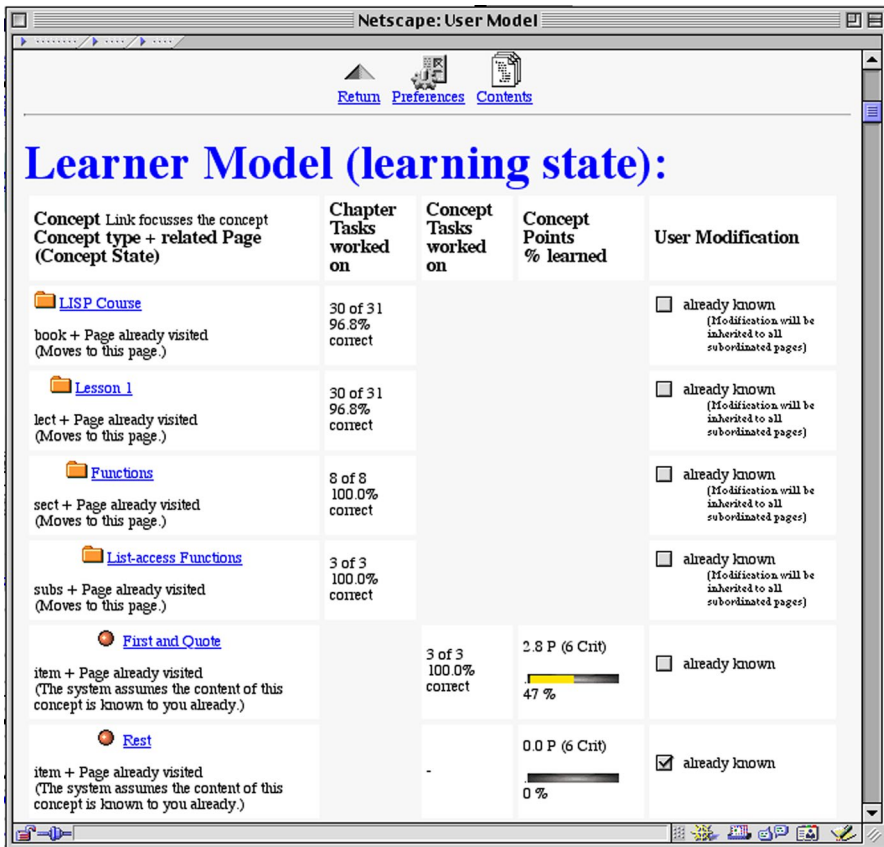


Fig. 1 Open and editable learner model in ELM-ART system (Weber & Brusilovsky, 2001)

component of intelligent tutoring systems (ITS). The original goal of open learner models is twofold. On the one hand, it attempted to show the learner what an AIED system thinks about her level of knowledge to make the actions of the system more understandable and the whole process more transparent. On the other hand, it allowed the learner to correct possible errors in the model, helping to mediate the imperfect learner modeling process. An example of simple learner control in OLM is the direct editing of the learner model (Weber & Brusilovsky, 2001). The ability to change the content of the OLM is sometimes stressed by calling this kind of model an *editable learner model*. In this case, it is the AI who determines the state of learner knowledge and displays it to the learner, while the learner has a chance to correct it by fixing obvious errors. Figure 1 shows the editable learner model in ELM-ART system (Weber & Brusilovsky, 2001). ELM ART is an online ITS to learn the programming language LISP and it uses the learner model of LISP knowledge to provide personalized question sequencing and navigation support. While different learners can start using the system with different levels of LISP knowledge, the system can only observe

user work within its borders and it could create poor personalization for learners who already know some aspects of LISP. Through the editable model, the learners can declare that specific concepts or groups of concepts are already known, and the system will take this into account in personalization.

An example of human-AI collaboration in learner modeling is a “negotiation” of the state of the learner model that a learner can do with a learner modeling agent (Bull & Pain, 1995). Here both parties could contribute to the final result rather than humans simply controlling the AI. Surprisingly, while it was one of the first approaches to engage learners in the modeling process, it received very little follow-up. Hopefully, the current interest in human-AI collaboration will encourage AIED researchers to return to the ideas of collaborative modeling.

Presentation and Selection

To make it clear what personalized content selection really means, I skip the step where the system selects relevant items with the help of the learner model and focus on the content presentation step first. Here the difference between no learner engagement, learner control, and human-AI collaboration is easier to demonstrate. Let us start with the “null” case where the presentation is fully controlled by the AI and the user has no say. This approach has been used in the very first AIED system Scholar (Carbonell, 1970) and was dominant for at least the first 25 years of research in the field. In this approach known as content sequencing, the AI content selection algorithm presents only one “next best” content item to work with, i.e., the next best problem to try (McArthur et al., 1988), the best example to review (Weber, 1995), etc. In this case, the learner has no choice but to accept the suggestion. In fact, in most sequencing interfaces, the learner may not even be aware that the selection of the next item was personalized - they may just think that is the way the system works for everybody. The classic sequencing approach was motivated by the expectation that “the AI knows best”, however, it did not always work well, since AI algorithms are rarely perfect and the state of the learner model could be frequently incorrect as explained above.

The analogy of “the AI knows best” approach in the “big AI” area are early recommender systems like WebWatcher (Joahims et al., 1997) which recommended one best link for the user to follow on the Web page or Google experiments with the “I feel lucky” search button that directly lead the user to a Web page that the algorithm considered most relevant. The problem with the imperfection of these approaches has been long recognized and both search and recommender systems switched to a much more flexible ranked list approach where the AI system presents a list of items ranked by the estimated relevance. This approach is a good example of human-AI collaboration in the presentation and selection of results. Here the AI does the work of careful selection and ranking, which is impossible for the user to do while the user does what AI cannot easily do, i.e., recognizing what the user really needs. In this collaboration, the user has the final word in selecting the most relevant content item while receiving help from AI through the presence of ranking. Essentially, the AI says “here is what you might like, and I think the things on the top you might like the most”. It is well known through many user studies that this

ranking “help” from AI users consider very seriously, but in no way do they always select the first item that AI considers as best (Keane et al., 2008). Ranking-based human-AI collaboration was a notable success and, despite criticism, survived intact through the first decades of research and application on recommender systems. Only in recent years has it started to lose ground to a more recent carousel-based approach to item presentation (Rahdari et al., 2022), which offered a more efficient approach to human-AI collaboration in the process of item selection. Not surprisingly, the majority of learner content recommender systems, and the currently most popular way of personalized content selection in AIED, adopted the ranking-based content presentation in its standard form (Drachler et al., 2015).

An alternative approach to human-AI collaboration in the process of item presentation and selection is *adaptive navigation support* (Brusilovsky, 2007). This approach was mostly pioneered by personalized education systems and was later adopted by many systems in “big AI”. Adaptive navigation support has been developed in the field of adaptive hypermedia (Brusilovsky, 2001) where learning content is presented as a form of hypertext or hypermedia. In this field, guiding users to the most relevant content means suggesting the best link to follow. Unlike WebWatcher (Joahims et al., 1997), which recommended one “best” link to follow, adaptive navigation support approaches in educational systems attempted to better engage learners’ own intelligence rather than relying on AI alone. Here AI still works in the background to decide which links are the best to follow, but AI advice is provided in a less direct form by adapting the presentation of links to user knowledge and goals using such techniques as link hiding (De Bra & Calvi, 1998), ranking (Papanikolaou et al., 2003), and annotation (Brusilovsky & Eklund, 1998; Weber & Brusilovsky, 2001). Here, as in the case of personalized ranking, AI and the user collaborate in finding the right content, however, AI typically provides better support than in the case of simple ranking by hinting why a specific item is relevant or not. A classic example of this support is “traffic light” link annotation used in such systems as ELM-ART (Weber & Brusilovsky, 2001) and InterBook (Brusilovsky & Eklund, 1998).

Figure 2 shows an example of using the traffic-light annotation in InterBook (Brusilovsky & Eklund, 1998). Here, the AI uses learner and content models to decide which connected pages are timely for the user to study and which are not, but instead of recommending the “next best” link in WebWatcher style, it adaptively annotates each link with a colored bullet where red color tells the user is not ready to work with the page while green color tells that it is just the right page to study.

Examples of more direct user control over the personalized presentation of information could be found in research on adaptive *presentation*, another popular technology in the area of adaptive hypermedia. Here a good example of “the AI knows best” case is provided by traditional AI-based approaches to adaptive presentation, which focus on generating content (for example, encyclopedia articles) that is adapted to user knowledge of the domain or other factors (Milosavljevic, 1997; Kobsa et al., 2001; Bontcheva, 2001). In this stream of research, the user is not able to control what is presented and is likely not aware that the presentation is adapted to her. Naturally, it could become a problem if the user model is incomplete or incorrect. For example, an AI could decide to *hide* some content from the user if it decides that this information is unnecessary or too complicated for the user to understand. Without

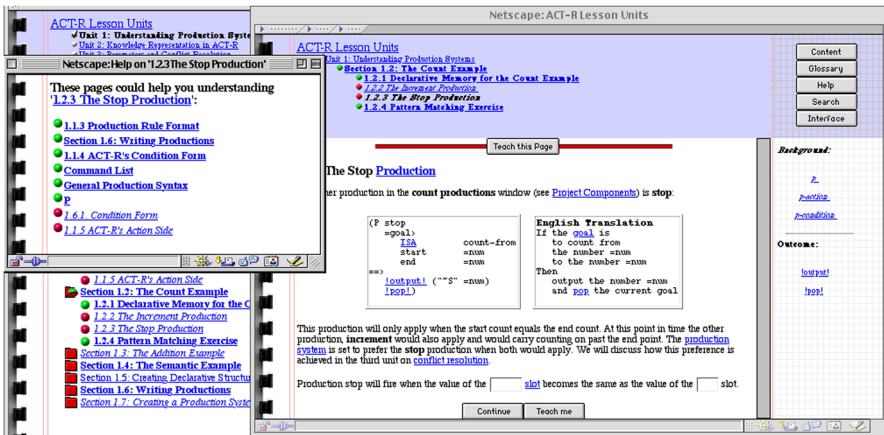


Fig. 2 Adaptive link annotation in InterBook (Brusilovsky & Eklund, 1998): AI-based relevance mechanism annotates links with colored icons indicating whether a specific link leads to a page that is just right, too simple, or too hard given the current state of the user knowledge. The control is left in the user's hands

any awareness and control over this personalization, the user might miss and never recover some vital information.

In contrast, the proponents of HCI-based approaches to adaptive presentation (Höök et al., 1996; Tsandilas & Schraefel, 2004) argued that user control over adaptive presentation is necessary for a usable system. This stream introduced several approaches that could be used to control adaptive presentation, such as adaptive stretchtext (Höök et al., 1996) and focus sliders (Tsandilas & Schraefel, 2004). Exploring the ideas of user-controlled adaptive presentation in adaptive educational systems, Czarkowski and Kay (2002) attempted to go from simple learner control to *scrutable personalization*. With this approach, a learner can *scrutinize* personalization, i.e., to see how the page is adapted to her and how this personalization decision was connected to the state of her learner model.

More recent research on user control and human-AI collaboration in content selection frequently combines features that were explored separately in earlier research. For example, a content recommendation interface in a personalized practice system for Database programming (Barria-Pineda et al., 2020), shown in Figure 3, combines elements of recommendation and adaptive navigation support as well as ideas of human-AI collaboration and scrutability. Here, AI selects several best database practice activities (examples, problems, or animations) given the learner's current level of knowledge and presents them in two ways, first, as a ranked list of the most appropriate activities and, second, by annotating links to the course topics and to learning activities (each link is displayed as a colored cell). Link annotations are made by placing stars of different sizes on the link cells, where the size indicates relevance (Figure 3A). AI also provides additional help in selecting the most appropriate activity by generating adaptive comments, which explain to the learner why this activity might be good for her current state of knowledge (Figure 3B). These explanations are provided when the learner mouses over a link. This example demonstrates that scrutability and

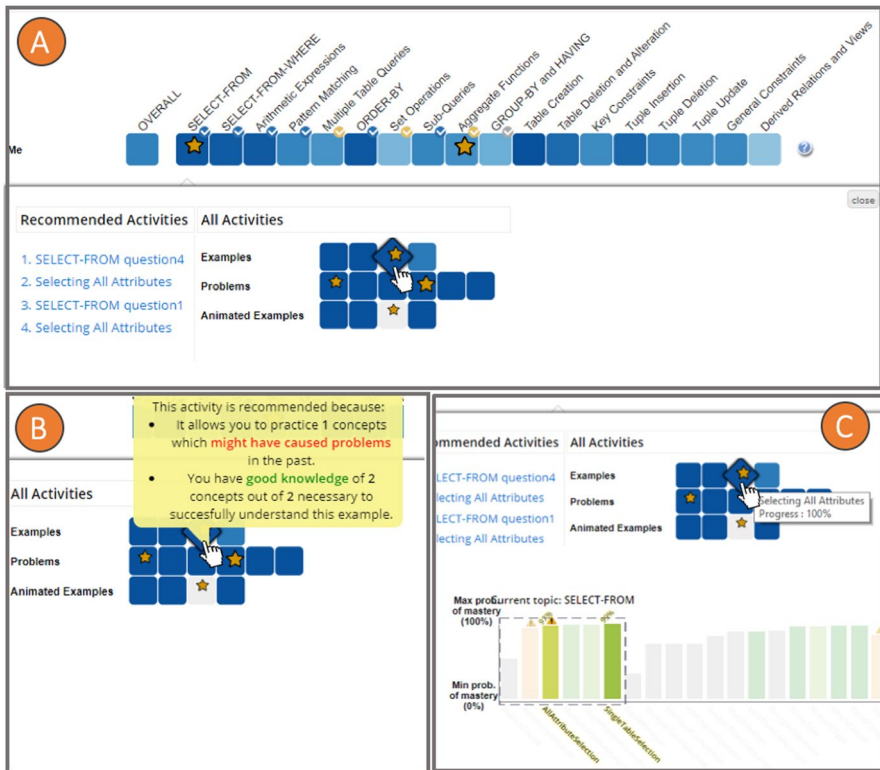


Fig. 3 Human-AI collaboration in content selection through a combination of ranking-based recommendation and adaptive link annotation (A). The system's explanation provides navigation support by explaining to the learner why an item was recommended (B)

explanations of AI decisions could make the process of learner control and human-AI collaborations more efficient by keeping the user better informed about the process.

The research on explainable AI (XAI) is rapidly growing in both “big AI” and AI Ed. A good review of XAI in AIEd could be found in (Khosravi et al., 2022). We will further discuss the need to make the AI side more understandable to the users in “[Transparency and Controllability](#)” section.

Determining What to Do Next

After discussing the user work with AI-selected items, it is time to return to the item selection process itself. Here, AI engages learner models, domain models, and content models to retrieve or generate content items that are most appropriate for the learner. After that, the items are typically ranked according to relevance and presented to the learner in some of the ways reviewed above. Surprisingly, neither in the field of AIEd nor in the “big AI” is the application of user control or

human-AI collaboration to this part of the personalized content selection process well explored. In other words, content selection in most cases is done by AI alone, while users are only engaged at the beginning of the process (learner modeling) or at the end (presentation and selection). To a considerable extent, the lack of research on user engagement in content selection could be the result of the complexity of modern content selection approaches, making it hard to control even for prepared users in regular recommender systems and even harder for less prepared learners in AIED systems (Czarkowski & Kay, 2002). However, there are several examples that demonstrate the opportunities for user control of the process.

The two main ways to control the selection process are to allow the user to choose one of the available algorithms to generate recommendations (Ekstrand et al., 2015) or to let the user control some parameters of the recommendation process. In the case of a content-based approach to the recommendation, the user could be allowed to express preferences about some content parameters of the desired items. For example, in an educational context, the user could be allowed to control the difficulty of selected items (Papoušek & Pelánek, 2017). In the case of a collaborative filtering approach to the recommendation, the user can control the group of peers used to generate it. For example, the PeerChooser system offered a graphical interface that allows users to enable or disable specific anonymous peers that are used to generate and rank recommendations (O'Donovan et al., 2008).

An analogy for the PeerChooser peer selection approach in the AIED field could be provided by a learner-controlled social comparison (Akhüseyinoğlu et al., 2022). In a regular social comparison approach implemented in the Mastery Grids system (Brusilovsky et al., 2016) and shown in the bottom part of Figure 4, the learner can compare her progress of knowledge (shown in green,

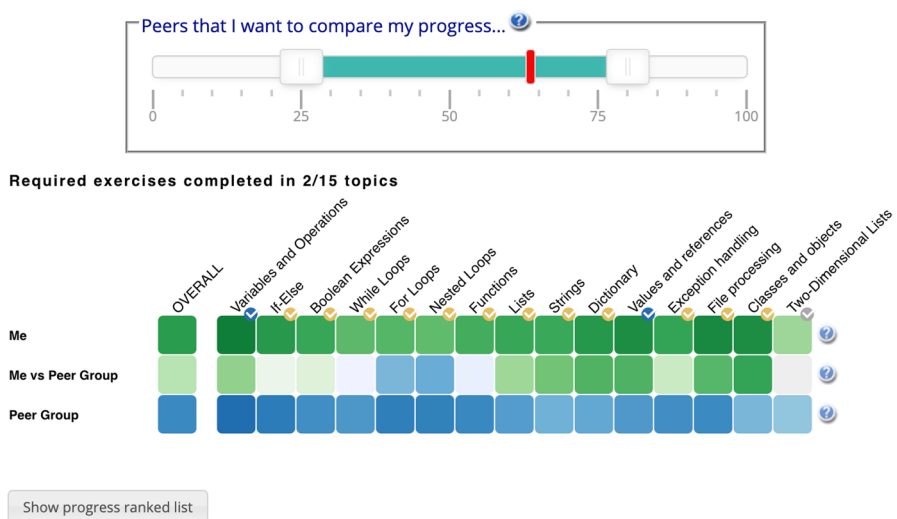


Fig. 4 User-controlled social comparison in Mastery Grids (Akhüseyinoğlu et al., 2022). A control panel with two sliders (top) allows the user to select the group of peers in the class who will be used to generate a social comparison of the learning progress (bottom)

topic by topic) with the progress of the class (shown in blue). It can help determine the topics where the learner is behind in class and guide her to the learning content in these topics as shown in Figure 3A. Using the whole class as a set of peers could be problematic, however. It could discourage strong learners while frustrating weaker learners, who will always see themselves well behind the class. To address this problem, the learner controlled social comparison mechanism shown at the top of Figure 4 allows learners to choose a subset of class members that they consider as “true peers” - for example, the very top of the class, the weakest learners, or the “middle” of the class as selected in Figure 4. Our studies show that the learning-controlled social comparison offers several benefits over the regular social comparison (Akhüseyinoğlu et al., 2022).

Transparency and Controllability

Although the three previous sections reviewed opportunities for learner control in all three stages of personalized content selection, the discussion cannot be completed without mentioning the issue of transparency. Whether the user is engaged in a straightforward learner control or a more complex human-AI collaboration, it is difficult for the user to provide a meaningful contribution without some understanding of the process. This aspect is known in “big AI” and recommender systems as *transparency*. Transparency is generally considered as the *other side* of controllability. On the one hand, as mentioned above, it is hard to execute control without transparency. On the other hand, full transparency cannot be achieved without some amount of user control where the user attempts to change various parameters and can observe how it impacts the results. Modern controllable recommender systems demonstrate many ways to make the recommendation process more transparent by visualizing some aspects of the process while offering the user some form of control over the process (O’Donovan et al., 2008; Knijnenburg et al., 2012; Ekstrand et al., 2015; Parra & Brusilovsky, 2015; Tsai & Brusilovsky, 2021).

To illustrate how transparency could be provided in an AIED system to better assist learners in the process of controlling and collaborating with AI, I want to show two examples related to the systems already discussed above. The first example shows how the content recommendation process explained above and shown in Figure 3 could be made more transparent. Here, the main source of knowledge for recommendation is the learner model, and the goal of the recommendation approach is to achieve a careful balance of new and well-learned concepts in recommended items. Consequently, good transparency can be provided by showing the state of the learner model (i.e., OLM) and highlighting which concepts are practiced in each recommended activity and what is the current knowledge state of these concepts. An example of visualization to achieve these goals is shown in Figure 5.

The second example shows how better transparency could be implemented for the case of user-controlled social comparison explained in the previous section and shown in Figure 4. Here, the comparative color annotation of course topics is determined by comparing the learner with the group of selected peers. To make the

| Recommended Activities | All Activities |
|--|--|
| <ol style="list-style-type: none">Swapping the First and Last Characters of a StringRepeating the Last Character of a StringAdding One String in the Middle of Another | <p>Examples </p> <p>Challenges </p> <p>Coding </p> |

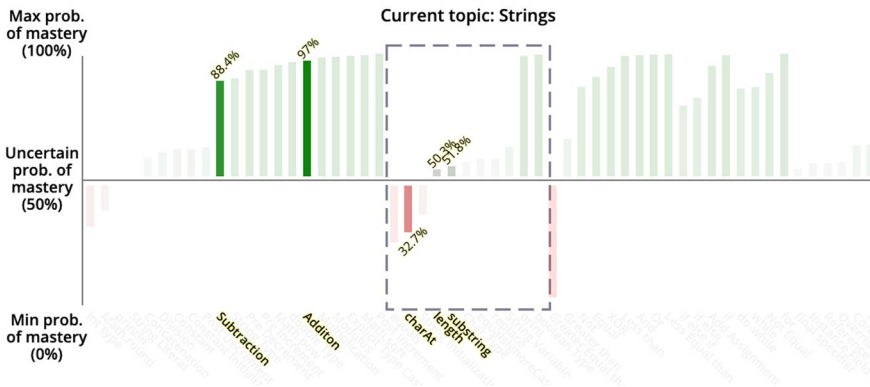


Fig. 5 Visualizing the state of learner knowledge through an open learner model and highlighting concepts that can be practiced by choosing an activity selected by the user (white square with the largest star) could make the next activity recommendation process more transparent (Barria-Pineda et al., 2021)

comparison of the user with her peer group more transparent, the system can offer an expanded view of social comparison showing the whole selected peer group and the position of the user in this group, as shown in Figure 6.




Fig. 6 A detailed list of peer learners showing their progress in knowledge and the position of the target learner in the list could offer some transparency to understand user-controlled social comparison (Akhüseyinoğlu et al., 2022)

Conclusion

In this paper, I have attempted to provide a case for learner control and user-AI collaboration in AIED. Although AIED systems led the work on user control through early work on open learner models, the field is now lagging behind the work on user control in “big AI”. By offering a range of examples for the implementation of learner control and user-AI collaboration in AIED, this paper hopes to encourage more research on this topic in AIED. In conclusion, it is important to stress that the work on user engagement in the work of AI in an educational system needs to be approached carefully and not by blindly following similar research in other areas of AI. Learners, especially younger learners, are a special category of users. They have much weaker knowledge of the domain and in many cases might not be ready to control or collaborate with AI. The age and preparation of the learner should always be considered when engineering learner control or collaborations with AI. However, many examples shown in this paper demonstrate useful and beneficial cases. I hope that the list of successful examples will expand further in the coming years.

References

- Ahn, J., Brusilovsky, P., & Han, S. (2015). Personalized search: Reconsidering the value of open user models. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (pp. 202–212). ACM.
- Akhüseyinoğlu, K., Milicevic, A. K., & Brusilovsky, P. (2022). Who are my peers? learner-controlled social comparison in a programming course. In I. Hilliger, P. J. Muñoz-Merino, T. De Laet, A. Ortega-Arranz, & T. Farrell (Eds.), *European Conference on Technology Enhanced Learning (EC-TEL 2022) Educating for a New Future: Making Sense of Technology-Enhanced Learning Adoption, volume 13450 of Lecture Notes in Computer Science* (pp. 1–15). Springer International Publishing.
- Bakalov, F., Meurs, M.-J., König-Ries, B., Sateli, B., Witte, R., Butler, G., & Tsang, Ä. (2013). An approach to controlling user models and personalization effects in recommender systems. In *Proceedings of the 18th International conference on Intelligent user interfaces, IUI '2013* (pp. 49–56). ACM Press.
- Barria-Pineda, J., Akhüseyinoğlu, K., Brusilovsky, P., Pollari-Malmi, K., Sirkiä, T., & Malmi, L. (2020). Personalized remedial recommendations for sql programming practice system. In *Workshop on Adaptation and Personalization in Computer Science Education at the 28th ACM Conference on User Modeling, Adaptation and Personalization, ACM UMAP 2020* (pp. 135–142).
- Barria-Pineda, J., Akhüseyinoğlu, K., Želem Čelap, S., Brusilovsky, P., Klasnja Milicevic, A., & Ivanovic, M. (2021). Explainable recommendations in a personalized programming practice system. In Roll, I., McNamara, D., Sosnovsky, S., Luckin, R., & Dimitrova, V. (Eds.), *22nd International Conference on Artificial Intelligence in Education, AIED 2021*.
- Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology*, *64*, 417–444.
- Bontcheva, K. (2001). Tailoring the content of dynamically generated explanations. In M. Bauer, P. J. Gmytrasiewicz, & J. Vassileva (Eds.), *8th International Conference on User Modeling, UM 2001, volume 2109 of Lecture Notes on Artificial Intelligence* (pp. 213–215). Springer-Verlag.
- Brusilovsky, P., & Millan, E. (2007). *User models for adaptive hypermedia and adaptive educational systems, volume 4321 of Lecture Notes in Computer Science* (pp. 3–53). Springer-Verlag.
- Brusilovsky, P. (2001). Adaptive hypermedia. *User Modeling and User Adapted Interaction*, *11*(1/2), 87–110.

- Brusilovsky, P. (2007). Adaptive navigation support. In P. Brusilovsky, A. Kobsa, & W. Neidl (Eds.), *The Adaptive Web: Methods and Strategies of Web Personalization, volume 4321 of Lecture Notes in Computer Science* (pp. 263–290). Springer-Verlag.
- Brusilovsky, P., & Eklund, J. (1998). A study of user-model based link annotation in educational hypermedia. *Journal of Universal Computer Science*, 4(4), 429–448.
- Brusilovsky, P., Somyurek, S., Guerra, J., Hosseini, R., Zadorozhny, V., & Durlach, P. J. (2016). Open social student modeling for personalized learning. *IEEE Transactions on Emerging Topics in Computing*, 4(3), 450–461.
- Bull, S. (1993). Towards user/system collaboration in developing a student model for intelligent computer assisted language learning. *ReCALL*, 5(8), 3–8.
- Bull, S. (2020). There are open learner models about! *IEEE Transactions on Learning Technologies*, 13(2), 425–448.
- Bull, S., Brna, P., & Pain, H. (1995). Extending the scope of the student model. *User Modeling and User-Adapted Interaction*, 6(1), 45–65.
- Bull, S., & Kay, J. (2007). Student models that invite the learner in: The smili() open learner modelling framework. *International Journal of Artificial Intelligence in Education*, 17(2), 89–120.
- Bull, S., & Pain, H. (1995). "Did I say what I think I said, and do you agree with me?" Inspecting and questioning the student model. In J. Greer (Ed.), *AI ED'95, 7th World Conference on Artificial Intelligence in Education* (pp. 501–508). AACE.
- Bunt, A., Carenini, G., & Conati, C. (2007). Adaptive content presentation for the web. In P. Brusilovsky, A. Kobsa, & W. Neidl (Eds.), *The Adaptive Web: Methods and Strategies of Web, volume 4321 of Lecture Notes in Computer Science* (pp. 409–432). Springer-Verlag.
- Carbonell, J. R. (1970). Ai in cai: An artificial intelligence approach to computer aided instruction. *IEEE Transactions on Man-Machine Systems, MMS-11*(4), 190–202.
- Cook, R., & Kay, J. (1994). The justified user model: A viewable, explained user model. In *Fourth International Conference on User Modeling* (pp. 145–150). MITRE.
- Corbett, A., McLaughlin, M., & Scarpinato, C. (2000). Modeling student knowledge: Cognitive tutors in high school and college. *User Modeling and User Adapted Interaction*, 10(2–3), 81–108.
- Czarkowski, M., & Kay, J. (2002). A scrutable adaptive hypertext. In De Bra, P., Brusilovsky, P., & Conejo, R. (Eds.), *Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'2002)*, volume 2347 of *Lecture Notes in Computer Science* (pp. 384–387).
- De Bra, P., & Calvi, L. (1998). 21670: A flexible adaptive hypertext courseware system. In K. Grønbaek, E. Mylonas, & F. M. Shipman (Eds.), *Ninth ACM International Hypertext Conference (Hypertext'98)* (pp. 283–284). ACM Press.
- Diessel, T., Lehmann, A., & Vassileva, J. (1994). Individualised course generation: A marriage between cal and ical. *Computers and Education*, 22(1/2), 57–64.
- Dimitrova, V., McCalla, G., & Bull, S. (2007). Open learner models: Future research directions. *International Journal of Artificial Intelligence in Education*, 17(3), 217–226.
- Drachsler, H., Verbert, K., Santos, O., & Manouselis, N. (2015). *Panorama of Recommender Systems to Support Learning* (pp. 421–451). Springer.
- Ekstrand, M. D., Kluver, D., Harper, F. M., & Konstan, J. A. (2015). Letting users choose recommender algorithms: An experimental study. In *Proceedings of the 9th ACM Conference on Recommender Systems* (pp. 11–18). ACM.
- Glowacka, D., Ruotsalo, T., Konuyshkova, K., Athukorala, K., Kaski, S., & Jacucci, G. (2013). Directing exploratory search: Reinforcement learning from user interactions with keywords. In *International Conference on Intelligent User Interfaces, IUI '2013* (pp. 117–127). ACM Press.
- Höök, K., Karlgren, J., Waern, A., Dahlbäck, N., Jansson, C. G., Karlgren, K., & Lemaire, B. (1996). A glass box approach to adaptive hypermedia. *User Modeling and User-Adapted Interaction*, 6(2–3), 157–184.
- Joahims, T., Freitag, D., & Mitchell, T. (1997). Webwatcher: A tour guide for the world wide web. In *15th International Joint Conference on Artificial Intelligence, IJCAI'97* (pp. 770–775).
- Kay, J. (1994). The um toolkit for cooperative user modelling. *User Modelling and User-Adapted Interaction*, 4(3), 149–196.
- Kay, J. (1997). Learner know thyself: Student models to give learner control and responsibility. In Z. Halim, T. Ottman, & Z. Razak (Eds.), *ICCE97, International Conference on Computers in Education* (pp. 18–26).
- Kay, J. (2001). Learner control. *User Modeling and User-Adapted Interaction*, 11(1–2), 111–127.

- Keane, M., O'Brien, M., & Smyth, B. (2008). Are people biased in their use of search engines? *Communications of the ACM*, 51(2), 49–52.
- Khosravi, H., Shum, S. B., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., & Gašević, D. (2022). Explainable artificial intelligence in education. *Computers and Education: Artificial Intelligence*, 3, 100074.
- Knijnenburg, B. P., Bostandjiev, S., O'Donovan, J., & Kobsa, A. (2012). Inspectability and control in social recommenders. In *Proceedings of the sixth ACM conference on Recommender systems* (pp. 43–50). ACM.
- Kobsa, A., Koenemann, J., & Pohl, W. (2001). Personalised hypermedia presentation techniques for improving online customer relationships. *The Knowledge Engineering Review*, 16(2), 111–155.
- McArthur, D., Stasz, C., Hotta, J., Peter, O., & Burdorf, C. (1988). Skill-oriented task sequencing in an intelligent tutor for basic algebra. *Instructional Science*, 17(4), 281–307.
- Milosavljevic, M. (1997). Augmenting the user's knowledge via comparison. In A. Jameson, C. Paris, & C. Tasso (Eds.), *6th International Conference on User Modeling, UM97* (pp. 119–130). SpringerWienNewYork.
- O'Donovan, J., Smyth, B., Gretarsson, B., Bostandjiev, S., & Hollerer, T. (2008). Peerchooser: visual interactive recommendation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1085–1088). ACM.
- Papanikolaou, K. A., Grigoriadou, M., Kornilakis, H., & Magoulas, G. D. (2003). Personalising the interaction in a web-based educational hypermedia system: The case of inspire. *User Modeling and User Adapted Interaction*, 13(3), 213–267.
- Papoušek, J., & Pelánek, R. (2017). Should we give learners control over item difficulty? In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization* (pp. 299–303). ACM.
- Parra, D., & Brusilovsky, P. (2015). User-controllable personalization: A case study with setfusion. *International Journal of Human-Computer Studies*, 78, 43–67.
- Rahdari, B., Kveton, B., & Brusilovsky, P. (2022). The magic of carousels: Single vs. multi-list recommender systems. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media* (pp. 166–174). ACM.
- Shneiderman, B. (2022). *Human-Centered AI*. Oxford University Press.
- Tsai, C.-H., & Brusilovsky, P. (2021). The effects of controllability and explainability in a social recommender system. *User Modeling and User-Adapted Interaction*, 31(4), 591–627.
- Tsandilas, T., & Schraefel, M. C. (2004). Usable adaptive hypermedia systems. *New Review in Hypermedia and Multimedia*, 10(1), 5.
- Waern, A. (2004). User involvement in automatic filtering - an experimental study. *User Modeling and User Adapted Interaction*, 14(201–237).
- Weber, G. (1995). Providing examples and individual reminders in an intelligent programming environment. In *AI-ED'95, 7th World Conference on Artificial Intelligence in Education* (pp. 477–484). AACE.
- Weber, G., & Brusilovsky, P. (2001). Elm-art: An adaptive versatile system for web based instruction. *International Journal of Artificial Intelligence in Education*, 12(4), 351–384.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.