

Selfit v2 – The evaluation of an Intelligent Tutoring System for Psychomotor Development

Laurentiu-Marian Neagu^{1,2}, Eric Rigaud¹, Vincent Guarnieri¹, Vladimir Ghita¹,
Sébastien Travadel¹, Mihai Dascalu^{2,4} and Razvan Rughinis^{2,4}

¹Centre of Research on Risks and Crisis Management, MINES Paris, PSL University
1 Rue Claude Daunesse, Sophia Antipolis, France

{eric.rigaud, vincent.guarnieri, sebastien.travadel}
@minesparis.psl.eu

²Computer Science Department, University Politehnica of Bucharest,
313 Splaiul Independentei, Bucharest, Romania

{laurentiu.neagu, mihai.dascalu, razvan.rughinis}@upb.ro

³Department of Management, University Politehnica of Bucharest,
313 Splaiul Independentei, Bucharest, Romania
vladimir.ghita@upb.ro

⁴Academy of Romanian Scientists, Str. Ilfov, Nr. 3, 050044, Bucharest, Romania

Abstract. An Intelligent Tutoring System (ITS) is a computer-based system that produces personalized tutoring through individualized, pedagogically sound, and easy-to-access educational material. Research groups explored various methods and assumptions for building efficient tutors in the cognitive field, with notable results in disciplines like physics, mathematics, and informatics. In contrast, the psychomotor domain is only lately exhibiting an intensive digitalization process. *Selfit v2* is a recently-developed ITS that aims to engage people in sports and improve the general health of the mass population. In this study, we assess *Selfit v2*'s utility and effectiveness in an experiment with forty-two users having low and medium training experience. The experiment used two adaptive strategies for tutoring – narrow and broad exploration spaces. *Selfit v2* evaluation showed promising results and highlighted the usefulness of ITS in the psychomotor field. The current work can be considered the foundation of a new crossroad between AI in education and psychomotor training, opening new research directions aiming to improve the population's general health through automated systems.

Keywords: Intelligent Tutoring System, Psychomotor Development, Personalization, *Selfit v2*, Usability Testing.

1. Introduction

The rise of information technologies, mainly driven by computer innovation, has revolutionized our interactions and learning processes [1]. The convergence of Internet science, Artificial Intelligence, and cognitive sciences have further created new tools in education, improving how educational content is produced and delivered while increasing education efficiency. Intelligent Tutoring Systems (ITSs) are Artificial Intelligence-based computer systems that provide adaptive educational experiences [2]. Starting with Carbonel in the 1970s [3], most approaches for building

ITSs have targeted cognitive skills, and multiple studies [4, 5] argued for the success of the employed methods in various STEM fields. However, the process of learning involves three domains: cognitive (knowledge-based), affective (attitudinal-based), and psychomotor (skills-based). Recent advancements in software engineering, computer vision, and sports training theory standardization [6] opened the opportunity for ITS application in the psychomotor skills field [7].

Selfit [8] is an ecosystem developed for training adult strength skills based on the traditional ITS architecture. The system supports strength profiling by considering fundamental strength movements with a performance scale from one to five for the upper body (push/pull horizontal/vertical) and the lower body (hip/knee dominant). It also supports strength workout prescription based on the initial evaluation - calibration challenges, trainee's readiness to train - before the training assessment, and the learning about the effects of past workouts - after the training assessment. A challenge for the tutor is determining the optimal sequence of activities to maximize the average competence level for all the targeted skills [9].

The ontology used for knowledge modeling within *Selfit* (i.e., OntoStrength [10]) is the first work, based on our literature findings, to describe the strength of psychomotor training. OntoStrength includes sports movements and movement patterns. OntoStrength also describes the psychomotor profile and plans for training programs for each phase - session, micro-cycle, mesocycle, and macro-cycle. Moreover, the ontology contains over 1,000 exercises labeled based on the involved muscles, movements, joints, materials required, videos, and corresponding difficulty.

The first version of *Selfit*, denoted as version 1.0 (i.e., v1), was tested within a simulated environment with populations of trainees and showed promising results for the adaptive tutoring component. *Selfit* v1 was also tested with a small group of trainees, exposing a few usability bugs. *Selfit* was further developed as v2 [11], and the improved system has been released recently. The new version included several bug fixes, performance improvements, and new features based on user feedback. Features include the ability to customize users' profiles (set up a profile picture, set a motto, birthdate, etc.), an option to pause and resume time-based exercises, an improved interface for inputting the number of repetitions per exercise, better tailoring of a session by inputting a list of materials available to train, added body areas trained statistics diagram, integrated Google Analytics, created a new protocol for pre- and post-tests, and others. This study aims to evaluate the effectiveness and usability of *Selfit* v2 in supporting individualized psychomotor training. To guide this evaluation, we address the following research questions:

- **RQ1:** To what extent does the adaptive tutoring mechanism (based on contextual multi-armed bandits) personalize exercise difficulty effectively, as perceived by users?
- **RQ2:** How does the exploration space (narrow versus broad) impact user engagement, training outcomes, and the system's learning performance?
- **RQ3:** What is the perceived user experience of *Selfit* v2 in terms of usability, motivation, and satisfaction?

Earlier publications introduced the *Selfit* ITS architecture [8; 11], the OntoStrength ontology for psychomotor knowledge modeling [10], and simulation-based evaluations of adaptive algorithms [8]. This paper presents a new contribution by evaluating *Selfit* v2 in a real-world longitudinal deployment with 42 users. We focus on analyzing the personalization efficacy of the contextual bandit algorithm, user engagement trends across exploration strategies, and the perceived user experience. No parts of the previous texts have been reused, and the present study builds upon those foundations to assess the system's performance in practical settings.

The current study presents lessons learned from an evaluation experiment performed with *Selfit* v2 to validate individualization algorithms' efficiency and

overall user experience. The paper is structured as follows. The next section presents the related work. The third section describes the experiment's aims and methodology. Then, the fourth section provides the outcomes of the training sessions and the results from the user experience feedback. The fifth section discusses the limitations of the system. Finally, the paper concludes with directions for future work through a list of envisioned applications.

2. Related Work

2.1 Intelligent Tutoring Systems in Cognitive Domains

Intelligent Tutoring Systems (ITSs) have traditionally focused on supporting cognitive skill development, especially in well-structured domains such as mathematics, physics, and computer science. Foundational works, including those by Carbonell [3] and VanLehn [4], have demonstrated the effectiveness of ITSs in mimicking human tutoring by adapting instruction to learners' needs. More recent systems employ probabilistic reasoning, rule-based modeling, and reinforcement learning to enhance personalization and engagement [2, 5]. These systems typically operate with clearly defined knowledge representations and learning objectives, where performance can be tracked and adjusted through real-time interaction.

Recent advancements in AI have significantly impacted intelligent tutoring systems and personalized learning. Lin et al. [12] emphasize the role of AI in creating individualized learning experiences, enhancing the effectiveness of ITSs. Similarly, Ilić et al. [13] provide a systematic review highlighting the positive effects of AI integration on student motivation and achievement in STEM education.

2.2 Emerging ITSs for Psychomotor Skill Development

In contrast to cognitive domains, the application of ITSs in psychomotor training remains underexplored. Psychomotor learning, which involves skill acquisition through physical activity and motor coordination, introduces unique challenges related to feedback, context variability, and physical effort tracking. Recent research has begun to address this gap, recognizing the increasing importance of digital tools in promoting physical activity and public health [7].

The application of AI in fitness has seen notable developments. Zhao et al. [14] explore how the perceived intelligence and anthropomorphism of AI fitness apps affect user engagement. However, few systems to date have tackled the complexity of modeling physical training progression or provided individualized guidance based on real-time performance and user state.

2.3 Knowledge Modeling and Ontologies in ITS

Structured knowledge representation is crucial for personalization in ITSs. Ontologies enable systems to reason over domain-specific concepts and learner states. In the psychomotor domain, *Selfit* introduces *OntoStrength* [10], a novel ontology that models strength exercises, movement patterns, training phases, and user profiles. To our knowledge, *OntoStrength* represents the first comprehensive effort to formalize psychomotor knowledge at this level of detail, enabling nuanced individualization based on muscle group, equipment availability, and user ability. This structured

approach distinguishes *Selfit* from other systems that rely on less formal or ad-hoc representations.

2.4 Adaptive Personalization via Reinforcement Learning

Adaptive recommendation in ITSs increasingly leverages reinforcement learning techniques, such as contextual multi-armed bandits, to optimize learning over time [9, 15]. These methods are well-suited for problems with uncertain outcomes and evolving learner profiles. Reinforcement learning's potential in healthcare is further explored by Bhattarai [16], who discusses its application in electronic health records for personalized medicine. Prior work has shown their promise in educational scenarios, where user engagement and performance can vary across contexts. *Selfit* applies this approach in a novel way by integrating user state (e.g., fatigue, motivation, available time and equipment) into its decision-making algorithm, dynamically selecting exercises to maximize user engagement and physical development. This use of contextual bandits in a real-world psychomotor setting is, to our knowledge, unique in current ITS research.

While significant advances have been made in ITSs for cognitive learning, systems targeting psychomotor development remain in early stages. Few leverage ontologies for structured knowledge representation, and fewer still apply adaptive algorithms that integrate real-time user context. *Selfit v2* bridges this gap by combining a psychomotor ontology with reinforcement learning to deliver personalized strength training plans. This paper contributes to the field by empirically evaluating the system's effectiveness in a real-world, longitudinal deployment, highlighting the promise of ITSs for health-oriented, psychomotor learning applications. This work is, to our knowledge, the first to evaluate such a system in real-world deployment, validating both the adaptive mechanism and user experience in a five-month study involving 42 participants, method which will be described further.

3. Method

The goal of an Intelligent Tutoring System is to provide more efficient teaching experiences to students [9]. This study's goal was to validate the software architecture and the interface of *Selfit v2* and to assess the extent to which the system learns the best load for each exercise so that learners are at their optimal difficulty level across a training program.

3.1 Participants

Forty-two trainees from France and Romania participated in our experiment between January and May 2022. Participants were recruited from close connections to the research groups behind *Selfit* and included software developers, students, teaching assistants, and sports athletes. Trainees used *Selfit* regularly and covered the most timeframes of the day (see Fig. 1, where the time shown was computed based on France's local time, namely GMT+2). The weekday with the most sessions was Tuesday, with three slots close to 80 training sessions: 5 PM, 6 PM, and 9 PM. Few or no sessions were performed during regular night hours (1 AM – 4/5 AM) and lunch (i.e., 12 AM - 1 PM).



Fig. 1. Training Sessions Count per day and hour slot.

3.2 Procedure

The trainee population was split into two groups based on the registration order within the system: group A (22) and group B (20). Both groups used the Contextual Multi-Armed Bandits ϵ -Greedy (0.1) algorithm, the best approach described in the initial simulations [8]. Multi-armed bandit addresses a problem in which a fixed limited set of resources must be allocated between choices to maximize the expected gain. The choice's properties are only partially known when picked, and it becomes better understood as time passes. The multi-armed bandit algorithm is a classic reinforcement learning problem that exemplifies the exploration-exploitation tradeoff problem. Reinforcement learning problems involve learning what to do and mapping a situation to the right actions to maximize a numerical reward signal [15]. The reward defined for our problem is the trainee's feedback related to the perceived difficulty of an exercise, in a particular state.

The first group (A) used a more expansive exploration space - the bandit arms for a movement area exercise, including all available levels filtered by user restrictions and preferences and the available materials. In contrast, the second group (B) used a narrow exploration space - the bandit arms for a movement area exercise included only the user's estimated level, filtered by restrictions and preferences, including available materials and muscles that are preferred to train. However, the second strategy may prove suboptimal considering the variability of user states (i.e., user context) – e.g., when the user is tired, more straightforward exercises may prove to be better; in contrast, when a user is in great shape, more strenuous exercises may be more efficient. An example of how the algorithm behaves for two users with the same training profiles using Tutor A and Tutor B after the calibration challenge can be seen in a simulated scenario in

Fig. 2.

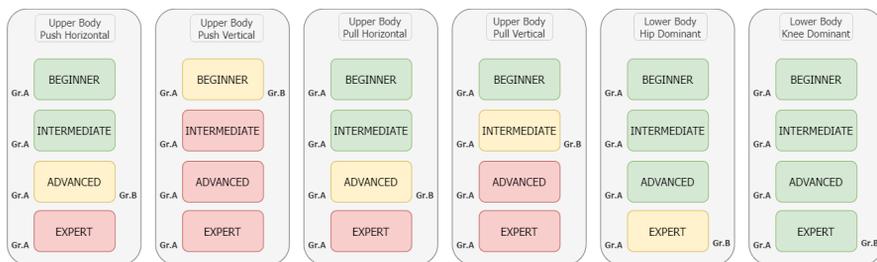


Fig. 2. Bandits Exploration Space per each group, based on the Calibration Challenge.

Even though the split per group was even, the number of exercises performed was unbalanced between the groups. The total number of exercises recommended by the bandit in this experiment phase was 1,521, from which Group A performed 1,297 (85.27%) and Group B performed 224 (14.73%). For the Upper-Body-Push-Horizontal movement type, the current level is Advanced. Therefore, exercises in Group A can be either Beginner, Intermediate, Advanced, or Expert, while exercises in Group B were only from the Advanced category.

Fig. 2 introduces the exercise counts performed in the most active training slot of the experiment – i.e., between 11th April - 23rd May 2022. The first part of a training session is the warmup, which lasts approximately 10 minutes. The warmup has different exercises based on the session's body target and includes, on average, 15 exercises to increase the athlete's body temperature and avoid the risk of injuries while training. Warmup exercises do not require feedback from the user. All activities are time-based and were not tracked in the experiment. An option to "Skip the warmup" was developed and released at the end of April to enable users to decrease the warmup duration when it was perceived as too long.

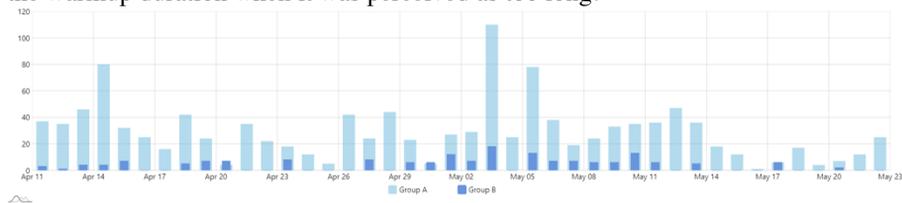


Fig. 3. Exercises count per training group between 11th April – 23rd May 2022.

The second part of training sessions is dedicated to blocs of exercises. The tutoring algorithm generated the content based on the session's target, the current training location (at home or gym), the available time (from 30 to 90 minutes), and the available materials. Moreover, users could indicate a list of muscles to train. The available options and corresponding mappings are presented in Table 1. Only content exercises were analyzed in the current experiment, which covered 315 training sessions.

Table 1. *Selfit v2* mapping of Time Availability and Count of Exercises Proposed.

Location	Time Availability (minutes)	# of Content Exercises
Home	30	5
	45	5
	60	6
Gym	60	6
	75	7
	90	8

The top 5 performers from Group A completed respectively: 30, 27, 25, 23, and 16, while the top 5 from Group B performers finished: 11, 6, 6, 5, and 5 sessions. Group A performed 254 training sessions, while Group B performed 61. Even though the group split was performed based on the random-like order of users joining the system, the statistics showed that the top 3 performers from Group A had trained more than all the trainees in Group B (20 users). Most of the time, users have chosen 30-minute sessions at home – 142 times, followed by 69 training sessions with a 45-minute length at home. Then, there are 62 sessions completed in a 60-minutes, both at home and at the gym. Thirty-five sessions (i.e., 11.11% of the total completed training

session) were conducted at the gym and within a 75-minute allocated time. Out of the 315, the remaining seven sessions were performed with a 90-minute time slot in a gym setup.

Fig. 4.a illustrates the statistics of training time chosen per session. The split between home and gym sessions was 225 training sessions performed at home (71.43%) and 90 training sessions at the gym (28.57%). Even though each session has a specific target, with a list of muscles to focus on, the users are requested to choose a list of complementary muscles to train. The list of available muscles was identical between home and gym training, and users could have picked as many muscles as they wanted or none. The experiment showed that the choice of users was diversified. Users frequently chose different muscles – see Fig. 4.b. Fourteen muscles were selected at least once by the users enrolled in the experiment. The top choices were glutes (157 times), abs (155 times), chest (122 times), back (114 times), and biceps (113 times). On the lower end, the choices were forearms (33 times) and neck (49 times), while the remaining were close to 100 selections. Before starting the training session, the system allows the definition of available training materials. If no material is selected, only bodyweight exercises are proposed. The material lists are different between home and gym training. Only elastic bands, kettlebells, and body weight are common in the two locations. Top users' choices of materials for home training were bodyweight (137 times), a chair (72 times), an elastic band (54 times), a kettlebell (41 times), and a table (39 times). When configuring a session in the gym environment, users select the dumbbell (84 times), barbell (80 times), pulley (63 times), body weight (58 times), or machine (57 times).

The choice of materials impacted the exercises for each session. The exercises chosen by the multi-arm bandit algorithm based on the session configuration and user profiles were in the vast majority bodyweight exercises – 1103 exercises, followed by exercises that required a dumbbell – 192 times, then an elastic band – 150 times, and a kettlebell – 131 times.

On the lower end, the wall was selected one time, the smith 6 times, and rings 15 times. The statistics showing the list of materials required by the user for practicing the recommended exercises can be seen in Fig. 4.c. *Selfit v2* recorded the immediate feedback for each exercise, represented by the Repetition in Reserve (RIR) per set and side (i.e., left or right) for unilateral exercises and the exercise level. The system also asked for user input on the overall perception of a session difficulty, called Rate of Perceived Effort (RPE), on a 10-Level scale, where value “1” is perceived as a tough session and “10” is very easy. Field RPE was not required. Out of the 315 sessions, no RPE was specified for 58 sessions. The most RPE provided was Level 8 (56 times), followed by Level 7 (48 times) and Level 6 (39 times), which means the sessions were mainly perceived as average in terms of difficulty. We considered values greater than eight as simple sessions, and values less than five as challenging; as such, we have 166 sessions (64.5% of RPE reported) marked between 5 and 8. On the other end, three sessions were voted with Level 1 (very hard), and two voted with Level 10 (very easy).

The statistics with levels for RPE and the number of associated sessions can be seen in Fig. 4.d. Each exercise described in the Domain model involved the movement coordination of one to four joints which, in some cases, is also linked with the exercise difficulty. The exercises proposed in the experiment had the following split, based on the number of joints: one joint (421 exercises), two joints (299 exercises), three joints (680 exercises), and four joints (79 times). Exercises are also categorized as unilateral and bilateral. This implies that the number of repetitions and sets per exercise should be executed separately on the left and the right (unilateral exercise) or simultaneously (bilateral). The current experiment included 1276 bilateral exercises and 203 unilateral exercises. Exercises were labeled on a 12-level scale and grouped into four categories: Beginner, Intermediate, Advanced, and Expert. These

categories also match the calibration challenges levels, which are helpful to group users per each movement type in the corresponding category. The statistics with the number of exercises recommended per each difficulty level are presented in Table 2. Most exercises recommended were at Level 3 (341), followed by Level 6 (240) and Level 5 (215). The least number of exercises were provided from Level 10 (9) and Level 7 (44). By categories, the ranking was the following – Beginner (636 exercises), Intermediate (512 exercises), Advanced (216 exercises), and Expert (115 exercises).

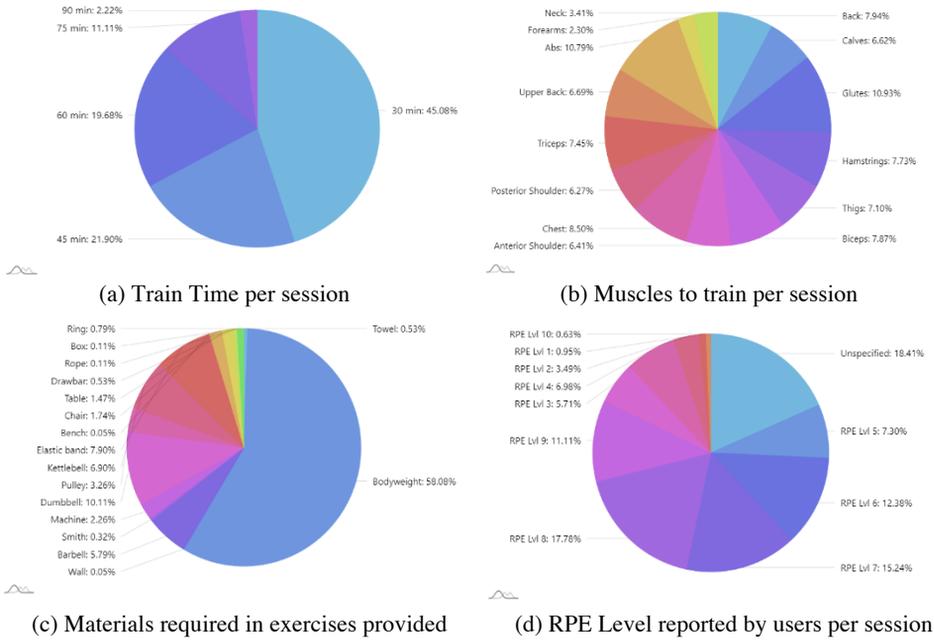


Fig. 4. Experiment Statistics on (a) Train Time, (b) Muscles to train, (c) Material used, and (d) RPE.

Table 2. Count of exercises recommended per exercise level.

Exercise Category	Exercise Level	# of Exercises	Exercise Category	Exercise Level	# of Exercises
Beginner	1	85	Advanced	7	44
	2	210		8	112
	3	341		9	60
Intermediate	4	57	Expert	10	9
	5	215		11	54
	6	240		12	52

4. Results

The experiment’s goal of learning was threefold. First, we aimed to validate the calibration challenges protocol and the classification of sports exercises. For this,

users have seen and interacted with while training. From those, users have filled in at least one value for RIR for 1442 of them.

Initially, *Selfit v2* allowed users to move throughout the training session without filling in any RIR value. However, a hotfix was released at the beginning of April, which made the RIR field mandatory. As a result, thirty-seven entries have no value for RIR and were excluded from the current analysis. For the significant sub-array, 21 entries had no RIR specified, and these entries were not used further. An exercise provided for a user at a certain time has between 3 to 8 RIR input values to fill in. Three inputs are the minimum – specific to 3 sets, bilateral exercise – see Fig. 6.a, while 8 is the maximum – 4 sets, unilateral exercise – required to input RIR for each part, left and right – see Fig. 6.b. RIR refers to the number of repetitions the user estimate he can additionally perform. If RIR is negative, the user performs less than required. If it is positive, the user can perform additional repetitions.

The absolute value of the RIR was considered in this stage to simplify the interpretation of the statistics. An optimal value for the RIR is viewed as a value that is close to zero. The RIR was computed both per each set performed and at the exercise level (sum of all sets in an exercise).

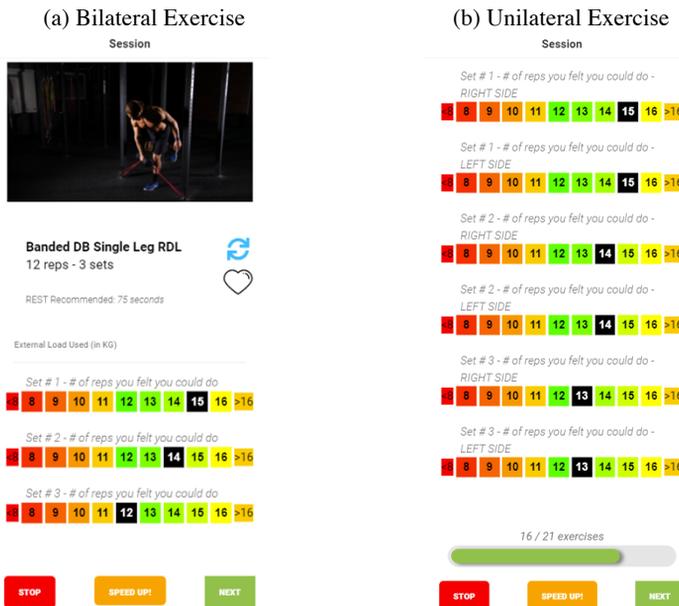


Fig. 6. User Interface for Inputting RIR per Exercise type.

The average RIR computed per set and per exercise for each experimental group can be seen in Table 3. The first line in the table covers the total amount of data from the experiment, interval [1,30]. It shows that Group B outperforms Group A on average RIR per set (34% smaller) and average RIR per exercise (26% smaller). The values are even more conclusive on the significant array [2,11], where Group B has an average RIR per set of 2.63, compared to 5.93 in Group A, and 8.08 is the average per exercise in Group B, while Group A is 14.84. The significant array is the most representative of the experiment, as the selected data is balanced between the two groups. The overall results obtained by analyzing the average RIR can be justified by

the fact that people who trained for 11 sessions or less in Group A have received easier or harder exercises. In the long run, after 12 sessions, this value starts to decrease. This is justified by the bandit algorithm which starts to provide better exercise recommendations as they learn from experience. This is discussed further in the next section. It can also be seen that while comparing lines one and two for Group A, the values decrease as the number of sessions per user increases. For up to 11 sessions, the average RIR per exercise is 14.84, while for all the data (up to 30 sessions), the average RIR per exercise is 11.34. A similar difference can also be seen in the average RIR per set.

Table 3. Average RIR per exercise and per set with different sessions interval sub-arrays.

# Sessions Interval Array	# Users / Exercises Group A	# Users / Exercises Group B	Mean RIR per exercise Group A	Mean RIR per exercise Group B	Mean RIR per set Group A	Mean RIR per set Group B
[1, 30]	22 / 1250	20 / 192	11.34	8.38	4.14	2.70
[2,11] (Significant Array)	8 / 121	11 / 165	14.84	8.08	5.93	2.63
[3, 8]	5 / 53	8 / 94	15.15	8.05	5.94	2.38
[5,12]	8 / 362	6 / 120	13.70	9.76	5.65	3.24
[3, 16]	14 / 688	9 / 151	12.89	8.65	5.04	2.80
[3, 30]	18 / 1218	9 / 151	11.27	8.65	4.13	2.80
[23, 30]	4 / 530	- / -	9.18	-	2.94	-

All other sub-arrays presented in Table 3 have shown a smaller average RIR over Group B than in Group A. For the users who performed fewer sessions, the difference is even more significant - see line 2 -array [2,11] -, and line 3 – array [3,8] - in the table. As the users trained more, the difference decreased. We can observe a difference between intervals [3,16] and [3,30] in Group A. In the first interval, the ones who have trained less have an average RIR per exercise of 12.89, compared to the second interval, where the value is 11.27. The average RIR per set is also 18% smaller for the second interval. For the trainees using the system consistently, it was shown that the average RIR in Group A tends to decrease over time and converges with the one in Group B. The last line of the table shows that for the four users who used the system for more than 23 sessions, the average RIR per exercise is 9.18, while the average RIR per set is 2.94, close to the values in Group B in fewer sessions. This analysis considers, at each subinterval (e.g., [2, 11], [3, 11]), all the user’s sessions without considering the current step in the training program.

The average RIR per session and exercise for users who trained more tends to decrease over time. To see if there was any improvement in RIR while training, a more granular analysis was performed per training microcycle - all the first microcycles for the users in Group A and Group B, all the second microcycles, and so on. The results for the first four microcycles per training group can be seen in Fig. 7. Fig. 7.a presents the first four weeks of training for Group A. The differences in RIR percentages between weeks are not significant, we see through a tendency for hard and soft fail RIR to decrease. Week 1 is 28.2%, while in week 4 is 21.7%. Fig. 7.b shows the RIR percentages in the first month of training for Group B. The hard failure RIR is decreasing after each week and becomes 0 starting in week 3, and this can be justified by the smaller space of exploration for Group B.

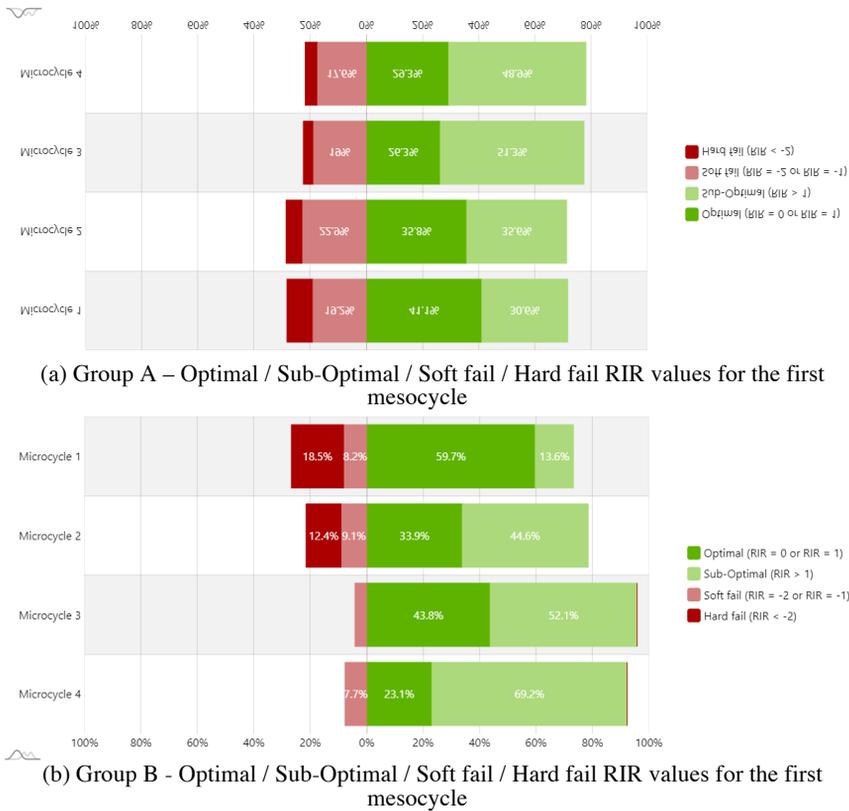


Fig. 7. RIR in the first mesocycle for (a) Group A and (b) Group B trainees.

Moreover, the overall cumulated optimal and sub-optimal RIR tends to grow as the users advance in the training program. As for comparing the differences in positive RIR between the two groups, we can see that, in each microcycle, the values in Group B surpass the values in Group A: microcycle 1 (A – 71.1%, B – 73.3%), microcycle 2 (A – 71.4%, B – 78.5%), microcycle 3 (A – 77.6%, B – 95.9%), microcycle 4 (A – 78.2%, B – 92.3%). We can conclude that Group B, with a narrow exploration space, provides better RIR on average than Group A in the first four weeks of training – the first mesocycle. It is also important to note that the Hard Fail was higher in Group B in the same timeframe. This is discussed further in the following sections. Group B had no trainee who followed the program for more than one mesocycle. In Group A, eight users performed more than four weeks of training (12 training sessions), and we observed that the percentages for optimal and sub-optimal RIR in the second microcycles tended to grow. For example, in Microcycle 6, the cumulated value was 87.1%, while in Microcycle 8, the cumulated value was 85.6%. Even though the values are not continuously increasing (Microcycles 5 and 7 have similar values to the ones in the first mesocycle for the same users), the bandit algorithms have a clear tendency to start gaining knowledge and to provide better recommendations. This is explored more in-depth in the following section. The acute effect at the exercise level is different between the two groups. On the one hand, the first mesocycle in the training program shows more optimal and sub-optimal RIR values on average in Group B than in Group A, which means the calibration tests are classifying the

trainees in the proper categories. If the classification were incorrect, more failures or less optimal values would have been shown in comparison with Group A. On the other hand, the labeling of exercises difficulty in *Selfit v2* is generally perceived as correct based on the results in Group B. Even though for a specific movement area level the user can receive exercises that get him to failure – see Microcycle 1 in Group B (26.7% of both soft and hard fail), this number of exercises is limited, and the majority in that specific level is labeled correctly. Based on all the above considerations, we observe high values for both optimal and sub-optimal RIR for Group B. We can conclude that the results obtained in the experiment validate the training testing protocol and the labeling of exercises' difficulty.

4.2 Results of Bandits' Learning

The second goal of the current experiment on learning was to validate the efficacy of the Contextual Multi-Armed Bandit algorithms in sports training planning. Top trainers from both groups were analyzed in this phase, as they had the most predictions to assess the learning. Group A, characterized by a vast exploration space from the bandit, required more states than Group B to provide a better reward across time. Therefore, users who performed more than 12 sessions were considered relevant in Group A based on the simulations presented by Neagu et al. [8]. The selected interval included 12 users and had from 57 to 169 actions taken by the bandit based on the current trainee state.

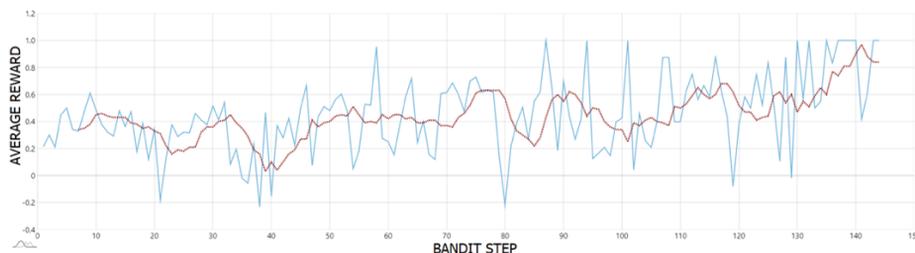


Fig. 8. Average Rewards and Moving Average (window size = 6 steps \approx 1 session) Rewards per Bandit Step for Top Performers in Group A ($>$ 11).

The results of the average rewards for top performers in Group A can be seen in Fig. 8. The figure presents each bandit prediction step, from 1 to 144, and the average reward across users tends to increase over time. Even though the most active trainee in Group A trained for 169 steps, we considered the averages of at least three trainees at each data point in the current analysis. Therefore, the number of bandit steps in Fig. 8 tops at 144, with the last step when there were three users with predictions. The information with all the top trainees in Group A, their corresponding bandit steps, and the number of sessions can be observed in Table 4. At each User ID, the number of trainees at that bandit step equals the value of that identifier (value of User ID). For example, in step 153, 2 trainees were using the system, whereas only one remained from step 154. Similarly for bandit step 78, 8 trainees were using the system, and from step 79 to step 81, there were seven trainees, and so on. The number of sessions is also displayed in the table, and the values are not always in descending order, as they are for the bandit steps. This is due to the length of specific sessions. For example, User ID #1, with 27 sessions, had longer sessions than User ID #2, with 30 sessions. Longer sessions use training session templates with more exercises, so the number of agent predictions is bigger.

Each session generates, on average, 6 agent actions. **Fig. 8** also shows the moving average on the dashed dark red line, which used a window of 6. The general increase of the reward over time is also easier to be observed with the moving average at the average session length window.

Table 4. Group A Top Trainers (>=12 sessions) Bandit Actions and Sessions Performed.

Top User ID	# Bandit Steps	# Sessions	Top User ID	# Bandit Steps	# Sessions
#1	169	27	#7	81	13
#2	153	30	#8	78	12
#3	144	23	#9	73	14
#4	135	25	#10	70	14
#5	93	16	#11	70	12
#6	84	12	#12	57	12

Tutor A shows promising results for the trainees who followed the training program for more than 12 sessions. The average values per step and the moving average show an overall increase, and to understand if there is any learning from the bandit, we take a closer look at the average values in Group A. We assessed the occurrence of high rewards (values greater than 0.5) over each step and computed the moving average for this array (see Fig. 9).

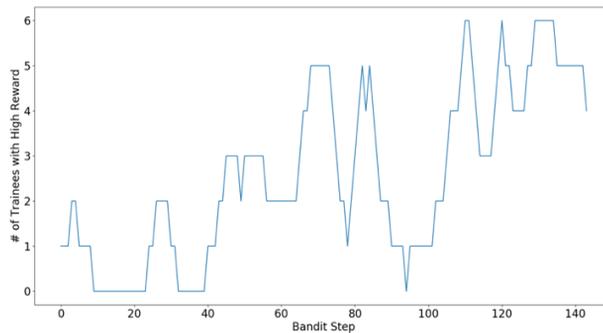


Fig. 9. Number of Trainees with High Reward (≥ 0.5) per Bandit Step.

In the first 50 bandit steps, the maximum reached was three users with a high reward. The value was obtained in three bandit steps (47, 48, and 49). For the second window, 50-100 bandit steps, the values are more excellent overall at most of the steps, and the maximum reached was five users with high reward, which was in 8 bandit steps (69, 70, 71, 72, 73, 74, 83, 85). For the last window, 100-144, we saw an overall increase in the high reward. The maximum number of users who reached high reward in this window was 6 in 9 bandit steps (111, 112, 121, 130, 131, 132, 133, 134, 135). The way the reward is computed makes the bandit agnostic as to whether an exercise is too easy or too hard (the reward considers only the absolute value). While comparing the RIR reported by the user per each bandit step, we observed that there was no learning of an exercise's positive or negative effect. We take a closer look at the predictions for a specific athlete and determine if the increasing rewards are linked directly with the learning. We use as an example User ID #3, for which we see an overall growth tendency of the reward. User ID #3 performed 23 training sessions, having a count of 144 exercises performed, for which the bandit made recommendations.

At the beginning of each training session, in each of the 23 performed, the athlete reported a good training shape. The user filled in a monitoring form before the session with values for fatigue restoration state, willingness to train, sleep quality, and stress level. Values were on a 10-points scale, where higher values denote that the parameter is not at risk. For example, for fatigue – a higher value means no fatigue, and for willingness to train – higher values mean the user is motivated to train, and so on. The average for these values is computed further to generate the current user state, between not good (when the average value is less or equal to three), good (when the average value is more significant than three and less or equal to seven), very good (when the average value is more significant than seven).

For example, User #3 reported an average value between 3 and 7 for each session, so the bandits classified him in a good state at each step. This means only one context state has been used for the current user. The median of the arms array for User #3 is 49. However, the overall split of exercise arms at a bandit step for participants in Group A, based on exercise difficulty, was most of the time equally balanced. The exercises were ranked in the four difficulty categories presented in Table 2. For the three bandit steps of User #3, there were 174 available arms. For the other three steps, there were 173 available arms. Moreover, for one step, there were 172 arms. On the lower end, one step had two available arms (step 65, see also Fig. 10) five available arms, one step had six available arms, and another one had seven available arms. Therefore, bandit choices on the 144 steps were the following: 103 steps used arms that were chosen only once (71.5%), 19 arms were chosen twice (26.3%), and only one arm was chosen three times (2.2%)– arm/exercise ID 480, which can also be seen in the first line in Fig. 10. Each line in Fig. 10 presents a step where the bandit must act. The list of available arms is shown on each line, and under each arm, the corresponding existing reward and the current step are shown, "-" value means that the arm was not selected yet. The selected arm is also displayed on each line, with a green background, and the reward obtained after executing that action is shown on the green line on the right. As the exploration factor was set to a low value, 10%, we can observe, e.g., for Step 11, the choice made by the bandit looks consistent with the expectations: only arm 480 has been previously explored, and the current reward is 1, so the bandit chooses again the arm 480, and the new reward obtained is 1. The same is true for Step 65, which is at line four in Fig. 10.

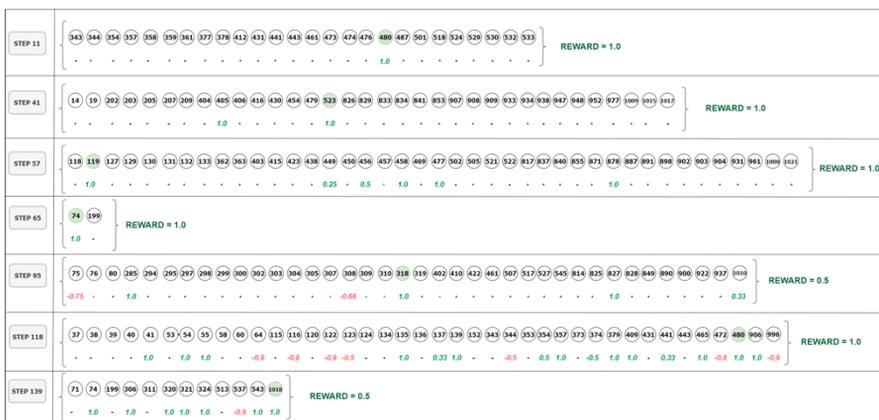


Fig. 10. Examples of Bandit Choices for User ID #3 based on the available arms

When there is more than one arm with the highest reward (value 1), the bandit chooses one of them randomly, as seen in Steps 41, 57, 95, 118, and 139 in the figure. When there was no history in any available arms, one value was picked randomly too. While plotting all the data, we also observed steps where the bandit chose to explore, even though high or average rewards were already available on some arms. Also, there were cases when the new reward was either positive or negative for existing arms. We can conclude that the bandit's choices impacted the overall increase of rewards for top trainers in Group A, so the bandits' learning is in line with our expectations.

These results support answering RQ1, indicating that the bandit-based tutor effectively adapted exercise difficulty over time, as reflected in improved RIR values and user-perceived training fit.

4.3 Results on User Experience

The experiment's third goal was to assess if *Selfit* ITS can adequately communicate with the learner in open environments. The exchange of information between the tutor and the learner is a critical component in the intelligent tutoring system field [1]. A user experience survey was conducted to assess whether the system is pleasant, easy to use, and efficient for the learners. This phase lasted for five months and involved 42 trainees from France and Romania. They enrolled in the system on different dates. The median value for the training session performed was five. *Selfit v2* did not store any personal data about the user. They enrolled within the system with a nickname, a password, and a security question, as a backup for forgetting the password. No real data about the user was required or available to fill in, such as email, phone number, name, or address. After logging into the system, users could configure their profile, fill in general information on their gender, birth year, and bio, or set up a profile picture. Out of the 42 users, only ten users filled in the optional information in their profiles. Only the bio and the profile picture fields were public and visible to the other members enrolled in the system. All the users were involved in the pre-test - the calibration challenge, performed when they started to train using *Selfit*. Ten also completed the post-test at the end of their training program. After the experiment, the trainees received a user experience survey to fill in. The structure of the survey was based on the AttrakDiff questionnaire. In addition, we added nine other questions to assess the training shape of the user, involvement in the experiment, and some open questions on the user's overall perception.

The AttrakDiff questionnaire [17] on a 7-point Likert scale was used to assess user experience while focusing on both pragmatic and hedonic qualities. The pragmatic quality of an object refers to how users perceive the system's potential "do-goal" to support them in achieving a goal (e.g., predictable, confusing, simple, complicated). The hedonist quality refers to how users perceive their "be-goal," namely *Selfit*'s potential to support pleasure in use and ownership (e.g., boring, interesting, novel, disappointing). Twenty-one trainees out of the initial 42 (50%) completed the questionnaire. Based on our knowledge, most of the trainees involved in the *Selfit V1* experiment [8] also used the new version and trained from January-May 2022. We aimed to assess if there is any visible improvement in the overall User Experience between the two versions, the main pain points, and strengths of the system, and if there is any link between the motivation to train and user experience. The AttrakDiff questionnaire values for *Selfit v2* can be seen in Table 5. The values in bold indicate a difference greater than 0.5 on the 7-point Likert scale between the two evaluations. Overall, user experience has improved in *Selfit v2*, as many hedonists and pragmatic quality values are better. On the other hand, the negative differences between the two

versions indicate that the users perceive the new system as more technical overall and feel like it does not fulfill their need to bring people together, the social component.

Table 5. *Selfit* v2 User Experience Feedback based on AttrakDiff questionnaire.

UX Quality	M (SD)	UX Quality	M (SD)	UX Quality	M (SD)
Pleasant	5.90 (1.26)	Connective	4.14 (1.64)	Human	3.85 (1.38)
Inventive	5.76 (1.71)	Simple	5.66 (1.65)	Professional	6.0 (0.83)
Attractive	5.19 (1.56)	Practical	5.66 (1.42)	Likeable	5.95 (1.21)
Straightforward	5.52 (1.36)	Stylish	4.90 (1.44)	Predictable	4.66 (1.49)
Premium	4.76 (1.54)	Integrating	4.95 (1.29)	Brings people closer	4.05 (1.30)
Novel	5.47 (1.49)	Motivating	5.90 (0.74)	Captivating	5.40 (1.39)

Of the 21 users who filled in the survey, 16 were male (76.2%), and five were female (23.8%). As per the age category - 19 users were between 20-40 years old, one was under 20, and one was over 40. The user profiles could have been more sportive. 20 users out of the 21 filled in their profession, from which 12 reported they were software developers, 4 were students, 1 was an interior designer, one university teaching assistant, one researcher, and one person working in economics. 11 users filled in the survey also their corresponding nicknames used in *Selfit* v2, to be identified with their corresponding survey data. Seven users were from Group A, while the remaining four were from Group B. They performed between 2 and 27 sessions, the mean was 9.63, and the standard deviation was 7.03. Users also had to rank the top features in the survey, up to 7 from a list of 17 features most asked for while training, features they want to see in the next version of *Selfit*. The most voted features can be seen in Table 6. Three features were voted on 13 times each by the users, most votes obtained per feature (61.9% of the total amount of users). Most users voted for "voice or sound guidance while training," which means, for the time-based exercises, the last seconds on the chronometer to also be announced with a sound, as the users do not follow their mobile phone screens all the time. Some users also proposed for this feature to be able to move within the app (next, pause, resume, previous buttons) using voice commands.

Table 6. Top-ranked features proposed by the users for the next version of *Selfit*.

Feature description	# votes	% users who want the feature
Voice or sound guidance while training	13	61.9 %
Auto Login / Remember me	13	61.9 %
Tips for each exercise / Extra description to understand better	13	61.9 %
Fix all the missing videos, improve some existing videos	12	57.1 %
Stretching module after training / Add Pilates, yoga exercises	10	47.6 %
Display more performance graphs / More statistics	10	47.6 %
Make warm-ups shorter, focus more on the core exercises	9	42.9 %

Another top-voted feature was "tips for each exercise / extra description to understand better." This feature was also reported during the experiment. Users have reported that more than the video for some exercises is needed to perform the exercise, and a complete description of body position and movements involved would be required. As for the third top-voted feature, "auto-login / remember me," we can state that this is more a missing implementation than a feature.

5. Discussion

We observed that participants from Group B quit the training much faster – no user performed the entire training program - 11 training sessions for the top trainer in this Group. In contrast, 12 participants from Group A followed the entire training program, and some top trainees – the top three performed 30, 27, and 25 sessions. The variability of the trainee state (low/high fatigue or low/high motivation) and bandit learning over time impacted the choice of exercises when generating a training session. Group A has demonstrated a higher motivation to train overall, which can be justified by the vast exploration space, which helped the bandit tailor the session with easier or harder exercises for each trainee's current context. There was a high impact of the exploration space parameter (*Selfit Tutor A* vs. *Selfit Tutor B*) in the choice of exercises and finally seen in the overall motivation to train due to the user's sensitivity to non-average exercises, based on his state. Our findings highlight the importance of the diversity of exercises through the broad exploration to maintain the motivation to train and keep the user engaged. The results for top trainees showed an increase in the average reward across time, and even though it did not follow the exact training planning for the simulation, the results of the experiments are slowly converging with the simulations, which was in line with our expectations. The Contextual MaB method for personalizing learning in psychomotor training has shown promising results overall.

Our work highlighted important limitations for both the current research process and for future researchers to consider when framing their studies. Next, we outline the identified limitations with corresponding descriptions.

The dataset obtained during data collection does not fully match the simulations' size. However, simulations performed in the virtual environment have demonstrated the learning of the contextual multi-armed bandits' algorithms. These showed that the users should train for at least 48 sessions to have a visible increase in the average reward across sessions, while in the actual experiment, the top performer trained 30 times. Due to the limited size of participants and time availability, we only tested some of the three bandits' algorithms used in simulations. Even though we could not create a larger dataset for this experiment, the results showed that the values converge to the simulated data. We also assessed several decisions taken by the bandits in different steps, proving that it is learning.

Even though we publicly released the system on mobile stores (i.e., Apple Store and Google Play) and the web, we had a limited budget to promote it and reach a high number of people. We promoted the experiment on our social platforms (LinkedIn, Instagram, and Facebook) and gathered 42 volunteers who enrolled in the program. The initial calibration challenge showed that most trainees were beginners in the sports training field. Very few of them were intermediate or advanced, participating in other sports activities. Their professions included software developers, scientific researchers, teaching assistants, and students. The majority was composed of males. Also, more than 75% of the users were between 20-40 years old. The sports individualization theory applies the same rules based on age differences and professions, considering the description of the trainee profile – gender, injuries profile, and diseases. Our findings support the generalizability of the system up to a certain degree. Younger (<20 years old) and older (>40 years old) are exposed to different injuries and diseases than the existing population. Our findings can be a springboard for future research to examine diverse populations.

In data collection while training, the measurement of perceived effort per exercise – repetitions in reserve (RIR) has been assessed using the subjective input from the trainee, which can be noisy. A more accurate measurement of the perceived effort is monitoring the heart rate while training and correlating with the input of RIR. The heart rate variability indicates if the user is training or not and what is the approximate

fatigue. It may also be a good indicator of counting the repetitions executed. Measuring the heart rate requires an external smart device, such as a smart band, smartwatch, or chest strap, which the participants involved in the experiment required while training. Also, integration with the third-party providers of such devices – Google, Apple, Fitbit, Garmin, Polar, and others. We implemented the integration with Garmin and Fitbit providers for heart rate monitoring. However, the users involved in the experiment did not have devices provided by these manufacturers to track data while training. We did not have the budget to provide such devices to the population in the experiment.

Another limitation is the complexity of the strength psychomotor training field, with multiple layers of knowledge representation, inner connections, and interconnections with other domains. Even though we obtained good modeling of the strongest field through *OntoStrength*, we did not map all the parameters that influence the individualization of sports training. This was due to time constraints and limited knowledge of cross-disciplinary expertise, such as the dietary dimension, where we needed additional knowledge.

Also, another limitation is generated by the overarching development of *Selfit* over the last four years, which was not heavily tested and may throw bugs in edge cases. Our team needed a Quality Assurance (QA) engineer assigned to develop test cases and execute them periodically. This exposes a vulnerability of the *Selfit* system overall. An overview of all available flows and edge cases should be described for any software to have a better view of the potential bugs which may occur. The analytics show a Code Maintainability Index of 74%, which is an average value for a code that is easy to support and change. More effort should be allocated to writing test cases, accompanied by code refactoring for better maintenance in the long term.

As outlined in RQ2, the exploration space significantly influenced outcomes: Group B's narrow space enabled faster convergence to suitable exercises, while Group A's broader space promoted longer-term motivation and higher engagement.

6. Conclusions and Envisioned Applications

Insufficient physical activity is a significant worldwide health concern affecting many people in developed and developing countries [18]. In addition, the changes in society, economy, and environment have resulted in low levels of physical activity and long periods of sitting [19], which can increase the risk of non-communicable diseases such as type 2 diabetes, cardiovascular disease, musculoskeletal disability, and certain types of cancer. The focus of the current work was to investigate the potential of intelligent tutoring systems (ITSs) in training psychomotor skills and to explore how these systems can be developed and utilized in a society increasingly concerned with the global public health [20]. *Selfit v2* is a recent release of an ITS for athletization designed to improve the general health of the mass population. The user's training constraints - limited time, fatigue, and volatile motivation- were integrated into a tutoring personalization approach based on RL contextual multi-armed bandits, shown previously to be efficient in educational scenarios. An experiment with real users was performed between January and June 2022, showing the proposed method's potential. However, further experimentation is necessary to fully explore the use of ITSs for training psychomotor skills. Addressing RQ3, user feedback revealed generally positive perceptions of *Selfit v2*'s usability and motivational quality, with actionable suggestions for further improving the interface and training guidance.

As digital health technologies become more integrated into daily life, there is an urgent need to move beyond one-size-fits-all fitness apps toward intelligent systems

that adapt to users' real-world constraints and psychomotor needs. *Selfit v2* addresses this gap by leveraging AI-driven personalization within a rigorous training ontology, targeting physical activity as a public health priority. By validating its use in real-life settings with a diverse population, our work not only offers a proof-of-concept for adaptive psychomotor tutoring but also lays the groundwork for scalable systems that can support health and well-being in increasingly sedentary societies.

Other systems targeting strength training can reuse the *OntoStrength* modeling and extend the work to other psychomotor fields – such as endurance or flexibility. The reinforcement learning implementation using contextual multi-armed bandits' algorithms was shown to be efficient for the adaptive personalization of learning activity sequences in the cognitive field [9]. The approach was proven suitable for scenarios when users have limited time and might face motivational issues when they are involved in long-term learning. The current paper used this approach for psychomotor training, also integrating the user's current state (trainee context), and it has shown promising results. The learning activities provided tend to maximize the student's skills over time. Our approach introduced a Python-based algorithm called RiERiT (Right Exercise at the Right Time), which estimates the trainee's learning progress and determines the appropriate exercise to be performed at a given time. This work can be used in other adaptive tutoring scenarios. Other systems that involve an adaptive component in psychomotor training may use the current approach as a starting point. The user interface model has been built as a monolithic system, using Microsoft tech stack – ASP.NET MVC. The code source will soon be published publicly on the GitLab platform. Specific features, such as calibration challenges, wearables integration with the most well-known providers (Garmin, Fitbit), session flow, and profile configuration by the users, may be reused by other researchers in their systems. Overall, governments are more and more interested in mass population health.

Moreover, the lack of physical activity is a significant concern for the more developed countries. Then, we envision the current work as the foundation of a larger-scale research project over the following years, focusing on digitizing the sports training field and developing several solutions that aim to improve the lifestyle of both beginners and intermediate trainees. The timing is right for such a research project, the technological advancements support the required developments, and our core team can provide the expertise to develop relevant solutions further.

CRediT author statement

Laurentiu-Marian Neagu: software, resources and writing – original draft. **Eric Rigaud:** methodology and writing – review and editing. **Vincent Guarnieri:** resources and data curation. **Vladimir Ghita:** software and formal analysis. **Sébastien Travadel:** conceptualization and supervision. **Mihai Dascalu:** writing – review and editing, and funding acquisition. **Razvan Rughinis:** project administration.

Acknowledgements. This research was supported by the project “Romanian Hub for Artificial Intelligence - HRIA”, Smart Growth, Digitization and Financial Instruments Program, 2021-2027, MySMIS no. 334906.

References

1. Woolf, B.P.: Building Intelligent Interactive Tutors: Student-centered Strategies for Revolutionizing e-Learning. Morgan Kaufmann Publishers, (2010)

2. Fenza, G.O., F.: Building pedagogical models by formal concepts analysis. Proc 13th Int. Conf. Intelligent Tutoring Systems (ITS), 9684, 144-153 (2016)
3. Carbonell, J.R.: AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction. IEEE Transactions on Man-Machine Systems, MMS-11(4), 190-202 (1970)
4. VanLehn, K.: The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educational Psychologist, 46, 197–221 (2011)
5. Butz, C.J., Hua, S., Maguire, R.B.: A Web-based Bayesian Intelligent Tutoring System for Computer Programming. Web Intelligence and Agent Systems, 4(1), 77–97 (2006)
6. Bompa, T., & Buzzichelli, C.: Periodization: Theory and Methodology of Training. Human Kinetics Publishers, Sixth édition (2017)
7. Neagu, L.-M., Rigaud, E., Travadel, S., Dascalu, M., Rughinis, R.-V.: Intelligent Tutoring Systems for Psychomotor Training – A Systematic Literature Review. In: 16h Int. Conf. on Intelligent Tutoring Systems (ITS 2020). Springer, Online (2020)
8. Neagu, L.-M., Rigaud, E., Guarnieri, V., Travadel, S., Dascalu, M.: Selfit – An Intelligent Tutoring System for Psychomotor Development. In 17th Int. Conf. on Intelligent Tutoring Systems (ITS 2021). Springer, Athens, Greece (Online) (2021) 282-286
9. Clement, B., Roy, D., Oudeyer, P.-Y., Lopes, M.: Multi-Armed Bandits for Intelligent Tutoring Systems. Journal of Educational Data Mining (JEDM), 7 (2015)
10. Neagu, L.-M., Rigaud, E., Guarnieri, V., Radu, E.I., Travadel, S., Dascalu, M., Rughinis, R.-V.: OntoStrength: an ontology for psychomotor strength development. Interaction Design and Architecture(s) Journal - IxD&A, 52, 101 - 118 (2022)
11. Neagu, L.-M., Rigaud, E., Guarnieri, V., Dascalu, M., Travadel, S.: Selfit v2–Challenges Encountered in Building a Psychomotor Intelligent Tutoring System. Crossley, S., Popescu, E. (eds) Intelligent Tutoring Systems. ITS 2022. Lecture Notes in Computer Science, 13284, 350–361 (2022)
12. Lin, C., Huang, A.Y.Q., Lu, O.H.T.: Artificial intelligence in intelligent tutoring systems toward sustainable education: a systematic review. Smart Learning Environments, 10(41) (2023)
13. Ilić, J., Ivanovic, M., Milicevic, A.: The impact of intelligent tutoring systems and artificial intelligence on students' motivation and achievement in STEM education: A systematic review. Journal of Educational Studies in Mathematics and Computer Science, 1, 5-18 (2024)
14. Du, Z., Wang, S., Deng, Z., Wang, F.: Unveiling the Power of AI Fitness Apps. Journal of Global Information Management, 33 (2025)
15. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press, 398 (2018)
16. Bhattarai, L.: Deep Learning and Reinforcement Learning in Electronic Health Records. Preprints, 2025041076 (2025)
17. Hassenzahl, M., Burmester, M., & Koller, F.: AttrakDiff: Ein Fragebogen zur Messung wahrgenommener hedonischer und pragmatischer Qualität. J. Ziegler & G. Szwillus (Eds.) Mensch & Computer 2003. Interaktion in Bewegung, 187–196 (2003)
18. Guthold, R., Stevens, G.A., Riley, L.M., Bull, F.C.: Worldwide trends in insufficient physical activity from 2001 to 2016: a pooled analysis of 358 population-based surveys with 1.9 million participants. Lancet Glob Health, 6(10) (2018)
19. Owen, N., Healy, G.N., Dempsey, P.C., Salmon, J., Timperio, A., Clark, B.K., Goode, A.D., Koorts, H., Ridgers, N.D., Hadgraft, N.T., Lambert, G.: Sedentary behavior and public health: integrating the evidence and identifying potential solutions. Annu. Rev. Public Health, 41 265-287 (2020)
20. McCuaig, L., Quennerstedt, M.: Health by stealth—exploring the sociocultural dimensions of salutogenesis for sport, health and physical education research. Sport, education and society, (2018)