Int J Artif Intell Educ (2025) 35:2117–2153
https://doi.org/10.1007/s40593-025-00463-z

ARTICLE

A Recommender System of Postgraduate Courses Based on Soft Skills: A Psychometric-inspired Approach

Luis Alberto Pinos Ullauri[1,2,3] · Alexis Lebis[3] · Abir Karami[4] · Mathieu Vermeulen[3] · Anthony Fleury[3] · Wim Van den Noortgate[1,2]

Accepted: 5 February 2025 / Published online: 20 February 2025

Abstract
Higher Education is constantly pushing to include soft skills in their curricula. An alternative could be to personalise their curricula so that students could better develop their soft skills. However, there are not many studies that investigate how to personalise the students' curricula based on their soft skills, especially if we consider multiple soft skills that need development. The aim of the article is to propose a recommender system framework based on soft skills in order to bridge the soft skills gap between the expected proficiency by employers and the actual proficiency of graduates. The approach is illustrated using real data from three cohorts of students that graduated in the years 2021, 2022 and 2023 at a French Higher Education Institution. We use a psychometric modelling approach to predict the soft skills proficiency of students within a genetic algorithm framework. We define three fitness functions and two aggregation methods, with which we can quantify the relevance of a set of courses across 10 different soft skills (e.g., Problem Solving, Leadership). The results show the recommendations to have, on average, a higher fitness than the actual courses taken by the students during the program. Moreover, there is significant evidence that the recommendations would allow the students to satisfy more of the soft skill targets compared to the courses the students actually took.

Keywords Recommender systems · Soft skills · Genetic algorithms · Courses · Fitness

Introduction

Soft skills are increasingly promoted not only by industry but also by Higher Education Institutions (HEIs). European Union politics induced HEIs to further emphasise soft skills in their program design (European Higher Education Area, 2016). For instance, the study by Cacciolatti et al. (2017) investigated the discrepancy between the industry's demand on both technical and soft skills alongside the universities' postures

Extended author information available on the last page of the article

on skills in the United Kingdom. This discrepancy resulted in a 'blame game' (Hurrell, 2016): a blame directed at individuals or at the educational system. Arvanitis et al. (2022) defined the soft skills gap as the misalignment between the required soft skills proficiency in work environments and the actual employees' proficiency. Moreover, Richa et al. (2021) argues this gap hinders the overall performance of workers, thus affecting their professional advancement as well as their organisations. This poses a problem, which may be addressed with further pedagogical inclusion of soft skills in HEIs.

In this context of the soft skills gap, it becomes important to understand and explain the soft skill proficiency of students during their studies. Pinos Ullauri et al. (2024) presented a psychometric approach to model the effect of postgraduate courses on soft skills. They argued that similarly to how students would differ in terms of soft skill proficiency (e.g., some students may be more proficient in Stress Management whereas others may be less adept in Leadership), courses potentially have different effects on soft skills. Estimates of the students' initial proficiency and the course effects allowed them to predict the students' soft skill proficiency after following a set of courses. Nevertheless, given that often the number of possible sets of courses is very large, knowing how to predict the proficiency after following a set of courses is not sufficient to recommend the most appropriate set. Thus, a next step could be to leverage this approach and complement it within a recommender system. This would allow students to identify better sets of courses, and therefore, better develop their soft skills and achieve a higher proficiency by the end of their studies. Consequently, this may help decrease the soft skills gap.

Recommender systems have been studied in education, such as the recommendation of learning items to students. Different approaches have been applied: Collaborative Filtering (CF) (Han et al., 2016), Content-Based Filtering (CBF) (Ibrahim et al., 2019) and other techniques such as Deep-Learning (Li & Kim, 2021). These types of approaches rely highly on considerable amounts of data from the learners, teachers, pedagogical items and their interactions. Therefore, these approaches generally leverage log data from Massive Open Online Courses (MOOCs) or Digital Learning Environments (e.g., the ratings learners have given to courses, the similarities between courses in knowledge content, times a learner spent on a task). Nonetheless, it can be quite difficult to collect sufficient data that are also relevant for predicting the effect on soft-skill proficiency for these approaches in HEIs, especially across entire academic programs. HEIs are very susceptible to policy changes from external parties such as political or international organisations. These changes may affect long-term data collection (courses may be added, removed or adapted). These constraints lead us to prioritise relative simple measures of proficiency, which can be feasibly collected over the span of a few years. Moreover, this data can comprise multiple cohorts, repeated proficiency measurements (e.g., once a year) and multiple soft skills (e.g., Problem-Solving, Stress Management). Additionally, the students' course history could also be collected.

Lebis et al. (2023) explored the issue of personalising students' curricula through simulation. They used another type of recommender system, often used in the field, based on genetic algorithms. Genetic algorithms commonly rely on a fitness function which determines how good a vector of variables can satisfy a problem. In their case, the fitness function from Lebis et al. (2023) determined the appropriateness of a sequence of courses for a particular student. They considered several constraints such as the number of credits students needed or course allocations. Similarly, the work by Shakhsi-Niaei and Abuei-Mehrizi (2020) framed long term course planning as an optimisation problem. They used aggregation in order to treat a multi-objective problem as a mono-objective case. They empirically chose the weights for each of the objectives, which dealt with constraints such as the preferences of students, the credits and difficulty of courses. Nevertheless, neither Lebis et al. (2023) nor Shakhsi-Niaei and Abuei-Mehrizi (2020) considered soft skill proficiency as the driving factor towards course recommendation. Moreover, if we consider various soft skills that require optimising, the problem becomes multi-objective, which can make the problem more complicated.

The complexity of recommending a set of courses is also another important subject to consider. The works by de Marcos et al. (2008) and Lebis et al. (2023) empirically tested the complexity of the recommendation problem. The complexity considerably increases with the number of courses. The higher the number of courses, the bigger the number of possible course sets that need to be tested. In principle, these may be solved by checking all the possible choices. However, the computational expense needed for an exhaustive search can be quite costly in practice, emphasising the need of a course recommender system for students.

Aim and Research Questions

In this study we explore the recommendation of HEI courses based on soft skills. The main aim of this article is to propose a recommender system framework, leveraging psychometric approaches in order to reduce the soft skills gap. The recommender system would find the most suitable set of courses that could allow the students to reach the desired soft skill proficiency at the end of their studies. In order to do so, the following research questions are devised:

RQ1: How can we recommend courses based on the students' soft skills proficiency and short-term small-size datasets from higher education?

RQ2: How can we represent and leverage both pedagogical postures and soft skills expectations in order to quantify the relevance of a set of courses across multiple soft skills?

The rest of the article is organised as follows. Section "Related Work" reviews related work on soft skills, educational recommender systems, soft skill proficiency regression models, and genetic algorithms. Then, Section "Methods" describes the data used in this study as well as our recommendation approach. We introduce the targeted soft skill profile and three different score functions. These are functions that deal with the comparison between the students' expected soft skill proficiency and their targeted profile at the end of the program. Additionally, we propose two aggregation methods that summarise the results from the score functions for each of the soft skills into a single fitness value. In order to illustrate the framework, real data from three cohorts of students from a French HEI is used, accounting for recommendations on 172 different course sets for students. Moreover, further specifications on the implementation of our recommender system are also described at the end of the section. Section "Results" presents the recommendation and fitness results. Finally, Section "Discussions and Limitations" presents a discussion of the advantages and limitations of the approach, as well as the conclusion of this article and possible strategies for future work.

Related Work

Before we delve into the data and methods, let us review key subjects from our proposed recommender framework such as soft skills, a psychometric model that can predict soft skill proficiency, recommender systems, and genetic algorithms.

Soft Skills

The definition of *soft skills* is a matter of debate among researchers (Hurrell et al., 2013). One source of the discussion stems from its relation with other terms such as competencies (Deist & Winterton, 2005), generic skills (Tuononen et al., 2022), 21st Century skills (Hu, 2024), graduate skills (Barrie, 2007), social skills (Notari et al., 2014), and employability skills (Succi & Canovi, 2019). In general, generic skills correspond usually to skills which can be leveraged on various situations whereas 21st Century skills emphasise the skills that are highly sought nowadays in the 21st Century. Employability skills focus mostly on the skills needed to apply to job positions. These jobs skills may not be limited to what is usually considered as soft skills, but also technical skills. Social skills are closely related to the interpersonal, rather than intrapersonal, aspect of soft skills.

In line with Almonte (2021), we consider soft skills as a group of non-technical, intrapersonal and interpersonal skills that we can use on various occasions. Taking that into account, a skill such as Decision Making could be considered as a soft skill since it is not technical in nature, it depends on personal and social factors (e.g., we may or may not have an easy time to arrive to a decision or we may or may not need to rely on someone else to make a decision), and it could be used in different situations (e.g., choosing the topic of an assignment).

There have been a few studies on the perception, application and modelling of skills in HEIs. Novais et al. (2023) proposed a fuzzy logic soft skill assessment framework during active learning lessons, based on self and peer assessments. In their ontology, they defined a 26 skills inventory grouped in 17 main soft skills. Tadjer et al. (2018) explored the students' perceptions on the importance of their soft skills. Tadjer et al. (2022) proposed a way to improve the soft skills mastery of students based on students' traces in problem-based learning environments. Sancho-Thomas et al. (2009) proposed a framework to promote the development and acquisition of teamwork skills across university programming courses. Muukkonen et al. (2022) investigated the students' perceptions on competence gains of 28 courses from two universities in Finland. They found the competence gain difference between courses to be statistically significant. This suggests that courses have different effects on the students' perceptions of competence development. The results by Pinos Ullauri et al. (2024) showed considerable variability of the effect of courses on soft skills. This further supports that not all courses help the students develop their soft skills in the same intensity. Moreover, given their generic nature, the soft skills proficiency could indeed be affected by any course.

Predicting Soft Skill Proficiency

If we were to predict the proficiency of various soft skills, a strategy could be to jointly model all outcomes using a multivariate regression model. Nevertheless, this approach can involve numerical challenges unless based on very strong assumptions about the association structure of the soft skills (Verbeke et al., 2014). Another, albeit simpler, approach could be to model separately each of the soft skills with the same regression model. In that way, we would have an ensemble of 10 models, one for each soft skill.

If the soft skills are measured with 4-point Likert scales, the use of ordinary regression may not be warranted. This is due to the discrete nature of the categories (i.e., the soft skills can be assessed with 1, 2, 3 or 4). Therefore, categorical regression can be used to model and predict the probabilities of the soft skill categories. Pinos Ullauri et al. (2024) modelled soft skills proficiency with multiple-membership ordinal logistic regression. Their model, shown in (1), predicts the natural logarithm of the odds of students being assessed with a category lower or equal than k (with $k = 1, 2, .., K-1$, and K being the number of categories). Similarly to any regression model, there is a linear combination of effects, albeit in the logit scale for categorical data. This linear combination depends on the students' course history, and their individual differences.

$$logit(P(sskill_{ist} \leq k)) = \alpha_{0ik} - (\beta_i N_{st} + \sum_{c=1}^{C} u_{ic} w_{cst} + \theta_{is}) \tag{1}$$

with i referring to the soft skill, s to the student, t to the stage, and c to the course, respectively.

The α_{0ik} are $K-1$ intercepts, which represent the log odds of an average student that has not taken a single course having a category k or smaller. The probability of the last category $P(\text{sskill} = K)$ is calculated by the complement, $1 - P(\text{sskill} \leq k-1))$. β_i is the fixed average course effect (same for all students) on soft skill i. A positive β_i would mean that by taking an average course, the odds of $P(\text{sskill} \leq k)$ decreases. This also means the odds of being assessed with a higher soft skill category than k increases. A negative β_i would increase the odds of being assessed with a lower category, thereby decreasing the probability of higher categories. N_{st} the number of courses that were followed by student s and stage t. In order to consider individual differences across courses as well as students, we could include regression coefficients for each of these. However, the amount of data needed to fit such a model would be incredibly large, especially considering the number of courses HEIs can offer as well as the number of students. Therefore, a strategy to account for these individual effects, without having to estimate each of them, is to define them as random effects. To account for the differences between students' initial soft skill levels, random student effects ($\theta_{is} \sim N(0, \sigma^2_{\theta_{is}})$) are considered. Similarly, courses may affect differently the students' proficiency. Therefore, u_{ic} represent the random course deviations from the overall fixed average course effect β_i (with $u_{ic} \sim N(0, \sigma_{ic}^2)$), whereas w_{cst} represent the course indicators. w_{cst} is equal to 1 if student s has taken course c by stage t, otherwise it is equal to 0. In this way, the individual course effect of a course c on soft skill i could be described by the average course effect β_i summed with its course deviation ($\beta_i + u_{ic}$). Moreover, empirical Bayes estimates can be obtained for the student and course random effects.

Recommender Systems

Ricci et al. (2010) defined recommender systems as software tools with the objective of providing valuable suggestions to users. Common examples of such applications are item recommendations in Amazon, video recommendations in media platforms (e.g., Youtube, Netflix) and song recommendations in music platforms such as Spotify. Data-mining techniques, such as CF and CBF, are amongst the most popular approaches for recommender systems, gaining considerably more attention with the Netflix prize competition, where several approaches were proposed (Hallinan & Striphas, 2016).

Collaborative Filtering CF relies on identifying different groups of users with similar preferences. In this context, a cluster is a group of users who have rated similarly the same items. The core assumption is that given an item that has been previously preferred by user 1, it could be preferable to another user 2, if both users belong to the same cluster. In essence, we try to predict the rating the user 2 could give to an item, regardless of the item's characteristics, based on the historical rating user 1 gave to that item, and the similarity in terms of preferences between the users 1 and 2. One of the main advantages is that CF does not require contextual knowledge regarding the users or items, but only their interactions. However, the two main drawbacks are the amount of data needed for the similarity calculations and predictions. The second limitation is the cold-start problem, where if a new user is introduced, it would prove difficult to calculate the similarity measure of the new user with the already known clusters.

Content-Based Filtering CBF focuses on the contextual knowledge of items and the interactions between a specific user and item. Unlike CF, this approach does not require historic interaction data from users other than the user who is being recommended. CBF normally proposes items which are similar to items previously preferred by an user. This can sometimes lead to overspecialisation (i.e., the recommendation of closely related items). Similarly to CF, it also requires considerable interaction data as well as contextual features of the items. For more details regarding both CF and CBF, please see Guruge et al. (2021).

Recommender systems have been applied in Education to support either learners, teachers or both (Yacobson et al., 2024). One of the main motivations for educational recommender systems is to help the learners in their decision-making process regarding the learning activities they plan to follow. This can represent a very complex task for the students, especially if they need to plan several tasks in advance (Lebis et al., 2023). The review by Guruge et al. (2021) presents commonly used approaches, such as CF, CBF and hybrid approaches. Nonetheless, these approaches are dependent on considerable amounts of data. That is the main reason most of them are applied in non-formal education contexts, MOOCs, Digital Learning Environments (where various types of interactions can be leveraged from the log history to extract important information for the recommendations of learning tasks) (Son et al., 2021; Dwivedi et al., 2018; Nabizadeh et al., 2020). Nevertheless, in a more realistic HEI setting, across an entire curriculum, the collection of sufficient data for these approaches could prove to be difficult, at least on the long term due to the varying policy changes on HEIs. If the users are students, and the items, courses, the interactions between them would themselves be limited. For instance, a type of interaction could be whether or not the student took a particular course. Unfortunately, the number of courses a student can take is constrained, therefore limiting the amount of data. Moreover, this type of interaction is not necessarily directly related to soft skills, unless other complementary variables are used to analyse the interplay between courses and soft skills.

Genetic Algorithms

Genetic algorithms are the most popular branch of evolutionary algorithms given their speed and efficiency (Son et al., 2021). The genetic algorithms' joint use in hybrid approaches with other recommender system techniques is also popular (Bobadilla et al., 2011; Alhijawi & Kilani, 2020). Moreover, genetic algorithms have also been widely used in e-Learning (Al-Muhaideb & Menai, 2011) as well as with simulated data (Lebis et al., 2023; Pinos Ullauri et al., 2023). Genetic algorithms are stochastic optimisation algorithms that resemble the natural selection process, where within each generation the strongest individuals are chosen as the parents. These parents help to reproduce the next generation of offspring. After several generations, we consider the strongest individual amongst the possible candidates as an approximation to the best solution.

More precisely, the individuals are represented by their genes. The genes can be coded as a vector of values whose representation depends on the problem. Figure 1 depicts the overview of a classical genetic algorithm workflow. First, an initial generation is generated. A generation is an ensemble of individuals. Then, amongst these individuals, various parents for the next generation are chosen. The crossover operation works with the parents to produce offspring, and thus, a new generation of individuals is created. Mutation may occur after the crossover, rendering the next generation. This generation may be the last generation, or serve as the generation for another iteration if the stop criteria have not been met (e.g., fitness of the best individual has not passed

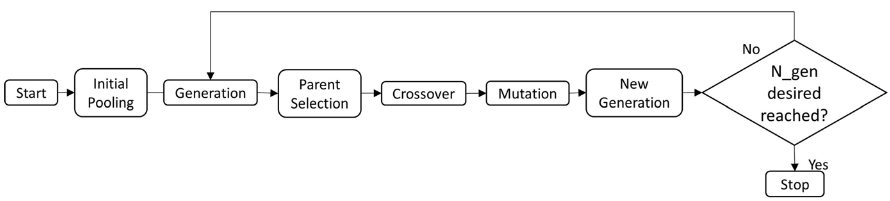

Fig. 1 General flowchart of Genetic Algorithms

a certain threshold) or a certain number of generations has been attained (e.g., 100 generations).

Methods

Data

The data used in this study comprises the soft skill proficiency and course history from three cohorts of students (totalling 884 students) that graduated from IMT Nord Europe in 2021, 2022 and 2023. The general engineering program is a 5-year program which integrates annual internships, of increasing length, in their curriculum. This particularity allows students to gradually become more familiar with industry. The students are assessed on their soft skills by internship tutors after each annual internship (at the end of each year, after they have taken their courses). A marking scheme, portrayed in the work by Pinos Ullauri et al. (2024), is used for the assessment across 10 different soft skills (please see Appendix A.1). The scheme describes a 4-point Likert scale of soft skill proficiency, where the students are assessed by their tutors, based on the description of their performance in tasks during the internships.

The soft skills considered in this study are listed in Table 1 alongside their identifiers for easier reading.

In the engineering program of IMT Nord Europe, the students essentially follow the same curriculum for the first 3 years. From the fourth and fifth years, the students must choose specialisations, where they have a catalogue of courses, which have no constraints other than expected technical skills. There are four specialisations: Energy and Environment (Specialisation 1), Industry and Services (Specialisation 2), Materials and Civil Engineering (Specialisation 3), and Digital (Specialisation 4). There are in total 104 courses in the dataset, which account for 100 specialisation courses and 4 transversal courses (that can be taken across all specialisations). From the specialisation courses 16, 18, 30 and 36 courses belong to specialisations 1, 2, 3 and 4, respectively. In addition, the students can take up to a total of 11 courses

Table 1 Soft skills studied in this article

Soft Skill	Soft skill identifier
Problem Solving	Soft Skill 1
Innovation and creativity	Soft Skill 2
Organisation	Soft Skill 3
Agility and adaptability	Soft Skill 4
Interaction and communication	Soft Skill 5
Project Management	Soft Skill 6
Conviction	Soft Skill 7
Stress Management	Soft Skill 8
Leadership and influence	Soft Skill 9
Decision making	Soft Skill 10

over the lasts two years. Since the courses can be freely chosen (from specialisation and transversal courses), the students' course history may very well diverge from one another. For instance, in Specialisation 1, the students have a pool of 16 specialisation and 4 transversal courses, which translates to $\binom{20}{11} = 167,960$ different combinations of course sets a student may take.

Table 2 shows an example of the entries of students 1 and 2 across the three last years, comprising the annual soft skills proficiency and course history. The first column corresponds to the student identifier whereas the stage column describes the year of the program the student was following at this measurement occasion (being the end of the 3rd year equal to zero, the end of year 4 equal to one and the end of year 5 two). The W_c columns are dummy variables, being 1 if student s followed course c, with $c = 1, 2, ..., 104$, 0 otherwise. The column labelled N describes the number of taken courses up until that stage, which can also be calculated as the row sum of the dummy variables. $sskill_{ist}$ is the soft skill level, assessed in a 4-point Likert scale by the internship tutor, of student s at stage t on soft skill i. For instance, the student with identifier 1 by stage 2, followed a total of $N_{s=1,t=2} = 11$ courses, amongst which were courses 1, 2 and 104, and soft skill 1 and 10 were assessed with a level of 3.

Training and Test sets Table 3 shows the number of students, whose soft skill proficiency data was assessed, per student cohort. If we had the complete data of the three cohorts and its 884 students, we would have $2,652 = (884\ students\ x\ 3\ years)$ data instances. Nonetheless, our dataset is highly sparse. It can be seen in Table 3 that there is no data from the 2021 cohort at stage 0. This is due to the fact that the soft skill assessment program had not yet started. Also, the last stage of the 2023 cohort is still missing, though its data is expected at the end of 2024. Moreover, there are students whose complete data cannot be accounted for. For instance, students who arrive from other HEIs in order to study the last two years (unknown soft skill proficiency at stage 0). Another example are students who drop out after stage 0 or 1. All of these particularities contributed to a total of 1,660 data instances we currently have in our dataset. Moreover, the number of data instances per course can be considerably low given that students can only take a maximum of 11 courses.

A commonly used rule of thumb for ordinary binary logistic regression argues there should be a minimum of 10 data instances per predictor in order to assure convergence and minimise overfitting (Peduzzi et al., 1996). Ideally, these data instances should

Table 2 Structure of student entries within the data set

Student	Stage	W_1	W_2	...	W_{104}	N	$sskill_1$	...	$sskill_{10}$
1	0	0	0	...	0	0	2	...	1
1	1	0	1	...	0	6	3	...	2
1	2	1	1	...	1	11	3	...	3
2	0	0	0	...	0	0	1	...	2
2	1	1	0	...	1	4	3	...	3
2	1	1	1	...	1	10	4	...	3
...	...	...	...	...	...	...	...	...	...

Table 3 Number of students per stage and cohort

Cohort	Stage 0	Stage 1	Stage 2
2021	–	250	141
2022	194	245	141
2023	339	366	–

be evenly distributed across predictors. However, as the number of possible courses is relatively high, and the amount of courses the students can take is limited, the data is highly imbalanced with courses with smaller amount of students compared to others. These limitations promoted us to consider the course effects as random effects rather than fixed effects, using the multiple membership approach.

Moreover, there are also other factors such as total sample size that can play an important role in the regression convergence (van Smeden et al., 2019; Riley et al., 2019). The dataset size as it is, is not sufficient for a successful convergence of the model parameter estimation procedure. Consequently, we used multiple imputation to generate various times the missing data from stage 0 of the 2021 cohort. We fitted the models on each of the imputed datasets and pooled the results, accounting for standard errors of the course effects due to the uncertainty of the imputed soft skill values. With this, the dataset size increases to 1910 data points, which is enough for convergence. As the data collection program is still active, we plan to update the dataset and its analyses each year.

Overall, considering the difficulties and complexities of the data, the entire dataset is used for the training part in order to mitigate convergence issues and obtain stable parameters. The test set is a subset of the training set, which consists of 172 students at stage 2 from cohorts 2021 and 2022. This accounts approximately for 20% of the students in the dataset.

Summary of the Course Recommender System Approach

Let us review and define certain terms in order to clarify the steps to follow in our proposed recommender system approach. First, we profit from the multiple membership ordinal logistic regression model, proposed by Pinos Ullauri et al. (2024), to explain the students' soft skills proficiency based on the course history and individual characteristics of the students. They defined *course effect* as the effect of a course on the students' proficiency at soft skill i. More precisely, if there are I soft skills, then each course would possess I effects, the same for all students. As a consequence, the trained models can predict the expected soft skills proficiency, given a set of courses. This means that by recommending new course sets, the expected soft skills proficiency can be calculated. In favour of simplicity, for the rest of the article the term *psychometric model* will be used to refer to this model. Second, in our recommender system approach, we require a goal to achieve. That is, a target for the students' soft skill proficiency at the end of the academic program. Since there are I soft skills, there are I targets. We propose the *targeted soft skill profile*, which comprises an ensemble of the targets for each soft skill i. Moreover, this profile may change depending on the program's specialisation m. Third, a *score function* has to be defined so that we can determine if a particular course set can help or not a student to achieve the target. This function assesses how the expected soft skill proficiency (given a particular course set) is compared to the target, providing a score. Since we may deal with various soft skills, we would need to aggregate and summarise all those scores. The result is called a *fitness value*. This value describes the relevance of a course set across all soft skills. The higher the value, the better the course set. Finally, we propose an algorithm, which essentially tests multiple course sets, and chooses a set of courses with the most optimal fitness value.

Figure 2 shows the overview of our recommender system approach. A set of courses serves as the main input for an ensemble of models $f_i([C_1, ..., C_n])$, which predict the soft skill proficiency $sskill_i$. The vertical green dotted line serves as a visual guide to delimit our contribution in this paper (right side of the dotted line) to the existing models (left side of the dotted line). It can be seen that the predicted proficiency $sskill_i, ..., sskill_I$ alongside the targeted proficiency $tarProf_{im}, ..., tarProf_{Im}$ (for specialisation m) are the inputs for the score functions sf. All of these I scores are aggregated into a single fitness value. Then, genetic algorithms are used to find the best course set, with the highest fitness value.

Targeted Soft Skill Profile

Now that we can estimate the soft skill proficiency of students, it is important to determine the desirable proficiency. Barrie (2006) found, through semi-structured interviews of academics, much variation within and between disciplines in their conception of the generic skills expected from graduates. Llorens Garcia et al. (2019) studied the skills gap in Spain in the Information Communication Technology (ICT) sector, and found the demand of soft skills for the ICT sector to comprise a subset of a larger portfolio of soft skills. In this study, they also identified the soft skills that employees usually master as well as the soft skills that they lack for the ICT sector. José-García et al. (2023) used Natural Language Processing and Text Mining techniques to identify and differentiate the required soft skills for various job position descriptions. Chamorro-Premuzic et al. (2010) defined a 15 soft-skills inventory and

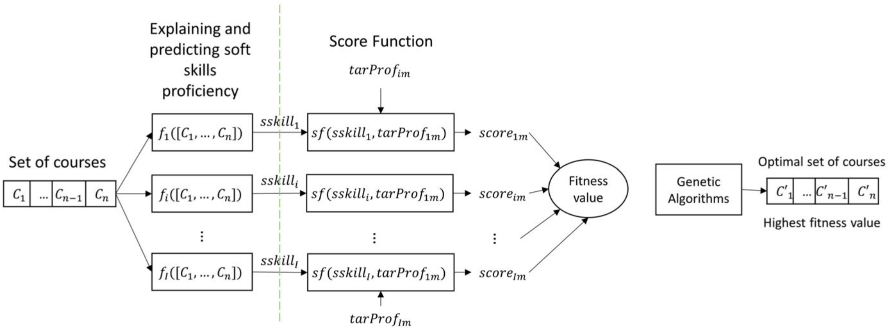

Fig. 2 Overview of the methodology of the proposed approach

found the demand of soft skills varied depending on the students' disciplines (particularly between science-related or social-related degrees). This suggests the demanded skill set of recently graduated students may differ not only in their technical skills, but also in their soft skills. In other words, in the same way as we would not expect the same technical skills from an aeronautical engineer compared to an industrial engineer, the soft skills of different specialisations could differ, thereby allowing distinct soft skill profiles for the different specialisations. In line with these findings, where the demand of soft skill proficiency may vary between degrees or specialisations, we propose the *targeted soft skill profile*. This profile would represent the ensemble of minimum expected proficiency targets for the soft skills for each specialisation.

Following the use of kiviat diagrams to represent skills and competencies in previous works (Schultz, 2010; Hanson et al., 2019), we present an example in Fig. 3. This is a fictitious example of how targeted profiles could differ between specialisations. In this example, the profiles are represented as polygons of 10 sides, depicting for each soft skill the targeted threshold. For instance, Specialisation 1 demands a rather balanced proficiency profile (2.5 out of 4 across all soft skills) whereas Specialisation 2 is mostly interested on the lower right part of the chart, but not as much in the upper left part of the chart. Specialisation 3, on the other hand, prioritises more the lower left part of the chart, demanding more proficiency in soft skills 6, 7, 8 and 9. Finally, Specialisation 4 has two prominent peaks in soft skill 3 and soft skill 7 and a lower expectancy on other soft skills.

Fitness

As we mentioned previously, the fitness value can be seen as the relevance of a set of courses regarding their compliance of the targeted soft skill profile. If we try to maximise this fitness, a course set with a higher fitness value is always preferred compared to another course set with lower fitness value. In order to calculate the fitness value, the following steps are followed. First, the psychometric model predicts

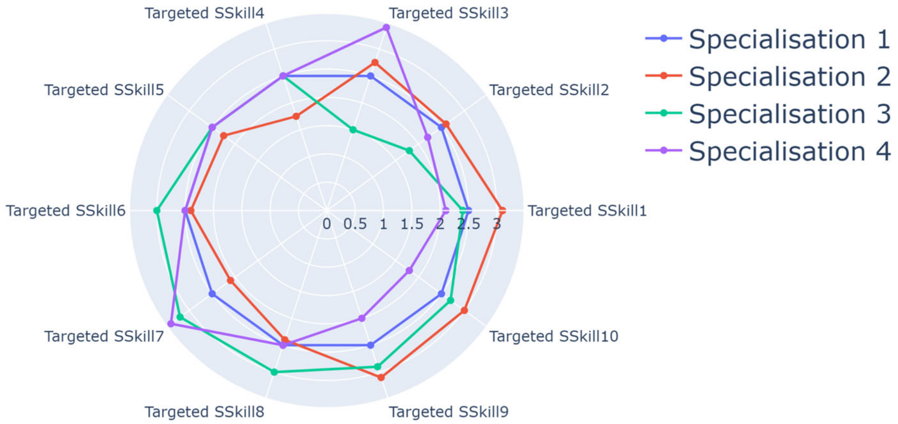

Fig. 3 Example of Targeted Profiles of 4 different specialisations

the probabilities on obtaining each of the possible categories for each soft skill, after having followed a specific set of courses. Second, based on the estimated probabilities, we calculate the expected soft skill proficiency. For instance, if the probabilities of a student s of being assessed with 1, 2, 3 and 4, respectively on skill i are equal to 0.05, 0.15, 0.45 and 0.35, then the expected proficiency $\overline{sskill_{is}} = 1(0.05) + 2(0.15) + 3(0.45) + 4(0.35) = 3.1$. Third, the expected proficiency $\overline{sskill_{1s}}, ..., \overline{sskill_{Is}}$ are compared to the targeted profile, yielding I scores (scaled so their range is within zero and one). However, these scores would be uni-dimensional (for each soft skill i, with $i = 1, 2, ..., I$). Hence, they would not summarise the goodness of a set of courses across all I soft skills.

This would normally turn our problem into a multi-objective one, where we would try to optimise each score simultaneously. Nevertheless, we propose an approach inspired from Multidimensional Item Response Theory (MIRT) (Adams et al., 1997; Reckase, 2009), which allow us to treat it as a mono-objective problem, and thus potentially diminishing the complexity of the optimisation problem. This is done by aggregating all the scores into an overall metric, the fitness value. Taking into account two classical pedagogical standpoints from MIRT, we propose two aggregation methods. Equation (2) portrays the first method, which is the average of the scores across all I skills. This aggregation allows compensation between surpluses or deficits of soft skill proficiency (a course set may not help the students reach the target on a soft skill, but it may help them outperform another soft skill target, making that on average, the student reaches the targeted level). This compensation is inspired by the compensatory MIRT models, where the latent abilities can compensate for each other.

$$CompFitness_s = \sum_{i=1}^{I} \frac{score_{is}}{I} \tag{2}$$

In contrast, a product allows from partial to no compensation if all scores are in the range between zero and one, compared to a sum because each score must be as high as possible to produce a higher fitness. Equation (3) shows the second aggregation method we propose, inspired by the partially or non-compensatory MIRT models.

$$PCompFitness_s = \prod_{i=1}^{I} (score_{is}) \tag{3}$$

$score_{is}$ refers to the soft skill gap of student s for soft skill i, this is the gap between the expected soft skill proficiency $\overline{sskill_{is}}$, and the targeted proficiency $tarProf_{im}$ of specialisation m, followed by student s, and soft skill i. The work by Richa et al. (2021) tried to identify the influential soft skills in the gap. They grounded their study on discrepancy theory of satisfaction and the theoretical underpinning of supply and demand of skilled talents (Farndale et al., 2010). In discrepancy theory (Tesch et al., 2003), if the skill proficiency is higher than what is expected, it is considered as a *positive discrepancy*, whereas if the proficiency is lower than the expected level, it is a considered a *negative discrepancy*. In accordance with these findings, we propose three different score functions that may deal with such discrepancies. All the functions are devised so that the proficiency is contained in the continuous interval between 1 and 4 and the scores between 0 and 1. This means there would be no negative scores. This would also entail that a proficiency of 1 is the lowest proficiency.

Figure 4 shows an example of how these functions would work on a particular soft skill, whose targeted threshold is at 3.0. Please note that these functions are examples, illustrating possible pedagogical behaviours or expectations, of how we could evaluate the discrepancy between the average soft skill proficiency and the targeted profile. Moreover, these functions could be mixed (or replaced) in order to fit the HEI needs. For instance, the HEI may be interested in strongly penalising the course sets that provide the largest negative discrepancies, while also wanting to provide a linear stance if the course sets allow the students a higher average proficiency than expected.

Baseline score function We consider a naive linear function that portrays direct proportionality between the expected soft skill proficiency and score. Moreover, the rate of this proportionality does not change if the expected proficiency is below or above the target. Figure 4 shows visually how all the score functions would work with a target proficiency level of three. The linear function is depicted in red. Please see Appendix B for further details on the mathematical grounding on all functions.

Logistic score function We consider a different function, which portrays a different behaviour than the baseline. In this function, small deviations from the target highly affect the scores, representing a rather strict policy (especially when using the non-compensatory aggregation). In contrast, larger deviations from the target are less penalised or prioritised (for both negative and positive discrepancies). Therefore, the recommended sets are expected to guarantee reaching the target. Figure 4 depicts the

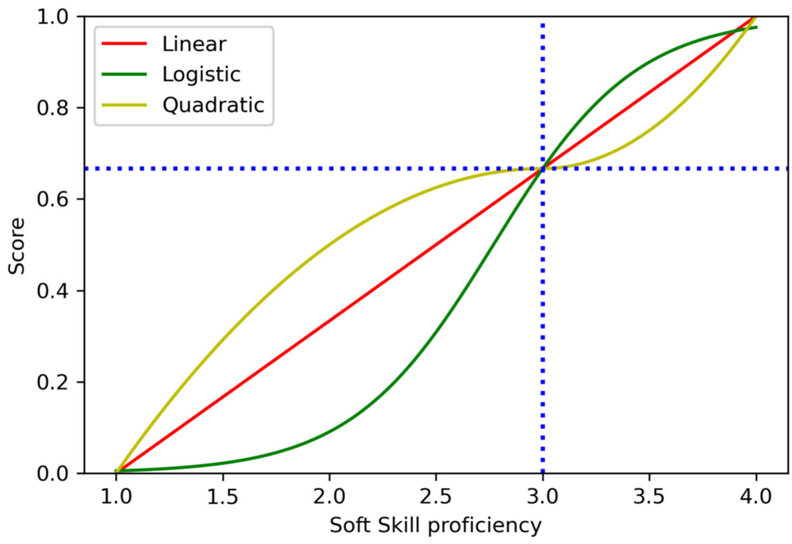

Fig. 4 Scoring functions. *Note:* The vertical line is located at the soft skill target. The horizontal line corresponds to the equivalent score from the linear function

logistic score function in green. Moreover, depending on an exponent we can adapt the steepness of the function.

Quadratic score function Finally, this function entails that small deviations from the target do not strongly affect the scores. This represents a more lenient judgement for the course sets that better allow students to achieve and surpass, as well for those that do not achieve the target at the end of the program. However, large deviations (both positive and negative) from the target do affect strongly the score. Figure 4 depicts the Quadratic score function in yellow.

Searching for the Optimal Course Set

Search algorithms are used to find a solution to a problem within the search space. The search space is an ensemble of all possible combinations of the variables (in our case, courses). Table 4 shows the number of possible course sets a student could choose depending on their specialisation over the last two years of the program. It can be seen that the search spaces vary considerably between Specialisation 1 and Specialisation 4 (from 167,960 up until more than two billion possible course sets). Using only the fitness function we could, in principle, test all possible course sets (known as Brute Force). However, this method is not viable for institutions with a large amount of freedom regarding the choice of courses. This is due to challenges in terms of computational processing for huge search spaces such as Specialisation 4, especially since we require inverse-logistic functions to calculate the probabilities for the four possible categories for each soft skill. Coding with Python in a computer with the following characteristics, Intel(R) Core(TM) i7-10810U CPU @ 1.10GHz 1.61 GHz and 32 GB of RAM, it took approximately 49.27 seconds to compute the fitness of 10000 course sets. This means that with a linear extrapolation, it would take approximately $2.31180144e9$ (49.27) $*e-4 = 1.33e10\,s \approx 131.83$ days non-stop to try all possible course sets of a single student of Specialisation 4 that took 11 courses out of 40 possible choices. Therefore, other techniques need to be favoured.

As we mentioned previously in Section "Related Work" and showed in the overview of the recommender system approach (see Fig. 2), we are using genetic algorithms in our implementation. Therefore, we start by defining how we represent the individuals in the population of candidate solutions.

Table 4 Possible course sets depending on the specialisation

Specialisation	Available courses	Max courses taken	Possible course sets
Specialisation 1	20	11	167960
Specialisation 2	22	11	705432
Specialisation 3	34	11	286097760
Specialisation 4	40	11	2311801440

Note: The recommendations are set at stage 2, with a maximum of 11 proposed courses

C_1	C_2	C_3	C_4	...	...	...	C_8	C_9	C_{10}	C_{11}
9	83	27	11	...	...	...	4	93	52	13

Fig. 5 Representation of a course set for a student in the genetic algorithm implementation. *Note: i* is the index of the course in the chromosome, and C_i corresponds to the identifier of the course in the catalogue $\in [1, 104]$

Individual Representation

Individuals are described by their genes. In our case, how the student would look like at the end of the educational program is described by the set of courses the student took. We will evaluate for different sets of courses how the student would look like at the end of the educational program, in order to recommend a set. The course identifiers are integers which start from 1 to 104 (since there are 104 courses). Moreover, the gene space depends on the specialisation and number of taken courses. An example of a course set of 11 courses is shown Fig. 5. In line with the psychometric model, the sequence is unordered, which means the course set does not impose a fixed order of the courses across time.

Initial Pooling and Parent Selection

For each student, the initial pool of individuals is obtained by drawing random samples of 11 courses from all the possible courses per specialisation (e.g., 40 courses for Specialisation 4). It is common to use a pool size of 100 individuals. This would produce 100 different fitness values, one for each individual. Figure 6 shows an example of a generation of 100 individuals. Next, from these individuals, parents are selected via a tournament, where various pairs of individuals *fight* against one another, having the course sets with higher values of fitness as the parents for the next generation of course sets. A common size for the parents is half the individuals, generally 50 for a pool of 100 individuals.

Crossover and Mutation

Crossover intents to combine parent course sets to create children course sets. We use uniform crossover, which deals with the mixing of the course sets by index. This means that from the courses that were chosen as the first course by both parents, one of the courses is copied to the child whereas the other one is not. This process is repeated

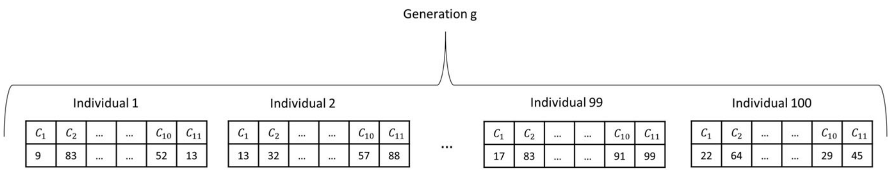

Fig. 6 Example of a generation of course sets

Parent 1	3	41	7	6	10	62	95	83	92	77
Parent 2	3	23	6	93	65	88	84	86	67	51
Child	3	23	7	93	10	88	84	83	67	77

Fig. 7 Example of the uniform crossover operation between two parents

iteratively for the other courses. Moreover, we implement a fallback mechanism, which checks for duplicates in uniform crossover. An example of how uniform crossover would work is also displayed in Fig. 7. The course set length is 10 courses, and the parents are displayed in blue and red. The crossover is performed by index. This means that for the fourth course there are 2 choices, course 6 and 93, respectively. For the third course, if course 6 would have been selected, the fourth course would have necessarily been course 93 to avoid taking the same course more than once.

The mutation process introduces randomness in the individuals changing some genes after a generation appears. We implement a single-point mutation, where one course is mutated in the child course set. That is, a new randomly generated course, from the gene space (removing the current courses from the course set from the possible choice to avoid duplicates), is put in its place. Hypothetically, this introduces a certain degree of variability without disturbing too much the convergence of the algorithm. An example of mutation can be seen in Fig. 8. The two algorithms are included in Appendix C.

Targeted Profile Selection

For simplicity, we chose per specialisation as targeted profiles, the students' average soft skills proficiency, at stage 2, across each of the 10 soft skills. Nonetheless, we acknowledge that the determination of such a profile is a study by itself and requires thorough debate of the stakeholders of the educational program (e.g., teachers and people from the labour market, alumni, current students, the HEI). Table 5 shows the targeted soft skill profile by specialisation.

Genetic Algorithm Hyper-parameters

The genetic algorithm hyper-parameters were chosen empirically, through various preliminary tests ensuring convergence, and are shown in Table 6. Given the stochastic nature of genetic algorithms, we perform 100 runs for each student for each combination of score function and aggregation method (Compensatory Linear, Compensatory

Before	3	102	19	26	87	95	66	8
After	3	102	19	26	93	95	66	8

Course id 87 is mutated to id 93

Fig. 8 Example of mutation of a given course set where a course is randomly changed

Table 5 Targeted profile selected for the genetic algorithm implementation

id	Soft Skill	Specialisation 1	Specialisation 2	Specialisation 3	Specialisation 4
1	Problem Solving	3.08	3.19	2.76	3.26
2	Innovation	2.83	3.00	2.70	2.89
3	Organisation	2.73	2.94	2.67	2.81
4	Adaptability	3.00	3.07	2.80	3.03
5	Communication	3.06	2.94	2.75	3.03
6	Project Management	2.40	2.56	2.33	2.23
7	Conviction	2.71	2.82	2.46	2.64
8	Stress Management	2.95	2.91	2.73	3.06
9	Leadership	2.59	2.75	2.38	2.64
10	Decision Making	2.86	3.02	2.73	2.99

The targets are the students' average soft skill proficiency (4-likert type scale, where the minimum is 1 and 4 is maximum)

Logistic, Compensatory Quadratic, Partially Compensatory Linear, Partially Compensatory Logistic, Partially Compensatory Quadratic) were performed. This means 600 genetic algorithm runs were executed for each of the 172 students of the test set.

Software

The algorithms were implemented in Python 3.9.12, with the use of the PyGAD Library (Gad, 2023), which is an open-source package to develop genetic algorithms. User-specified functions for the crossover, mutation and fitness calculation were included as well as several modifications and methods to work with our dataset. More information regarding the implementation can be found in Appendix C.

Hardware

A cluster, composed of 1,388 heterogeneous threads was used to execute our genetic algorithms. Computationally speaking, a cluster is an ensemble of various computers working as a single system. Clusters can run several calculations in multiple threads.

Table 6 Genetic algorithm hyper-parameters setup

Specialisation	Number of Generations	Population size	Crossover rate	Mutation rate	Seeds
Specialisation 1	300	100	0.80	0.25	[1,100]
Specialisation 2	300	100	0.80	0.25	[1,100]
Specialisation 3	500	100	0.80	0.50	[1,100]
Specialisation 4	500	100	0.80	0.50	[1,100]

Each run was performed separately, under a single thread and a different seed (e.g., Compensatory Linear from student with id 434 from Specialisation 3 with seed 20). Having such computational parallelisation option at our disposal, we were able to compute the equivalent of 866 days of computation (under a single thread at 2.67 GHz) into a few weeks of computation.

Results

Table 7 shows a summary of the results by Aggregation type, specialisation and score function. The fourth column shows the average fitness with its standard deviation, across students, of the courses taken by the students whereas the fifth column displays the same metrics for the recommended courses. The sixth column shows the average percentage increase in fitness of the recommendations compared to the actual taken courses. The seventh column depicts the average difference of soft skill satisfied targets between the recommendations and actual taken courses (e.g., if our recommendations allowed the achievement of 5 out of 10 soft skill thresholds, while the actual courses would allow only 3, the difference would be $5 - 3 = 2$). It can be seen that across all score functions and aggregation types, the average increase in the fitness is positive. This increase varies from 2.2% in Quadratic in Specialisation 1 to 21.9% in Logistic in Specialisation 4 in compensatory terms. The increase is considerably higher in partially compensatory going from 23.1% to 997.2%. The average difference between the recommendations and the actual courses that satisfied the targets is positive at all times. This difference ranges from 1.81 up until 3.76 in Compensatory and is also statistically significant for all cases. This means that, on average, the recommendations helped satisfy more targets than the actual courses the students took. Please note the average target difference does not explain by how much the targets are satisfied, but whether or not the targets were reached. Therefore, the fact that they are pretty similar amongst the various fitness functions does not pose a problem.

It is important to add that all students in the test set benefit from the recommendations with a higher fitness, compared to the fitness of the actual courses they took. This entails the overall fitness increase was always positive across fitness functions and students. Nonetheless, the same cannot be said about the uni-dimensional scores across each soft skill. Figure 9 shows a box-plot of the percentage increase (across students) of the fitness across all soft skills in Specialisation 3 for the Compensatory Linear Fitness. It can be seen that the increases vary across the 10 soft skills. Soft Skill 5 has both the largest median and the largest interquartile range (length of the box, which contains 50% of the students' percentage increases). In general, this means this is the most positively affected soft skill, though there is also considerable variability compared to the other soft skills. Soft Skill 4 and 6 are amongst the least positively affected soft skills, having a small part of negative increases. Nonetheless, the compensation across soft skills allows other surpluses to balance the overall proficiency, thereby allowing for flexible combinations of soft skill proficiency. Please see Appendix D for more results and detailed boxplots of all the fitness functions from each specialisation.

Table 7 Summary table of results

Aggr. Type	Spec.	Score Function	Avg Fit Real (SD)	Avg Fit Recomm (SD)	Avg Fit increase	Avg target diff
Comp	1	Linear	0.59 (0.04)	0.64 (0.04)	8.9%	3.36***
Comp	2	Linear	0.57 (0.03)	0.63 (0.04)	9.4%	2.26***
Comp	3	Linear	0.56 (0.02)	0.61 (0.04)	8.4%	2.28***
Comp	4	Linear	0.59 (0.04)	0.65 (0.04)	9.7%	3.48***
Comp	1	Logistic	0.57 (0.08)	0.67 (0.07)	17.9%	3.43***
Comp	2	Logistic	0.49 (0.06)	0.60 (0.08)	**21.9%**	2.26***
Comp	3	Logistic	0.57 (0.05)	0.66 (0.07)	15.4%	2.31***
Comp	4	Logistic	0.56 (0.08)	0.66 (0.08)	19.3%	3.69***
Comp	1	Quadratic	0.61 (0.01)	0.62 (0.01)	**2.2%**	**3.76*** **
Comp	2	Quadratic	0.63 (0.01)	0.64 (0.01)	2.2%	2.26***
Comp	3	Quadratic	0.55 (0.01)	0.56 (0.01)	2.6%	**1.81*** **
Comp	4	Quadratic	0.62 (0.01)	0.63 (0.01)	2.7%	2.69***
P. Comp	1	Linear	5.7e-3 (3.7e-3)	1.2e-2 (6.5e-3)	140.8%	3.31***
P. Comp	2	Linear	3.7e-3 (1.9e-3)	9.8e-3 (5.9e-3)	162.8%	2.26***
P. Comp	3	Linear	2.8e-3 (1.1e-3)	6.9e-3 (3.8e-3)	132.2%	2.31***
P. Comp	4	Linear	5.7e-3 (3.8e-3)	1.4e-2 (8.8e-3)	158.9%	3.65***
P. Comp	1	Logistic	6.1e-3 (7.6e-3)	2.2e-2 (1.9e-2)	594.0%	3.45***
P. Comp	2	Logistic	1.1e-3 (1.3e-3)	8.7e-3 (9.6e-3)	**997.2%**	2.26***
P. Comp	3	Logistic	3.9e-3 (2.9e-3)	2.0e-2 (1.6e-2)	390.0%	2.32***
P. Comp	4	Logistic	5.7e-3 (7.3e-3)	2.5e-2 (2.5e-2)	709.6%	3.68***
P. Comp	1	Quadratic	6.2e-3 (7.7e-4)	7.6e-3 (1.1e-3)	23.5%	3.79***
P. Comp	2	Quadratic	9.2e-3 (1.2e-3)	1.2e-2 (2.1e-3)	**23.1%**	2.26***

Table 7 continued

Aggr. Type	Spec.	Score Function	Avg Fit Real (SD)	Avg Fit Recomm (SD)	Avg Fit increase	Avg target diff
P. Comp	3	Quadratic	2.3e-3 (1.6e-4)	3.0e-3 (5.4e-4)	27.0%	1.90***
P. Comp	4	Quadratic	7.0e-3 (7.7e-4)	9.0e-3 (1.7e-3)	28.8%	2.71***

Notes: Aggr. Type stands for Aggregation type, Comp for Compensatory and P. Comp for Partially Compensatory, Spec. for Specialisation, Avg. Fit Real for Average Fitness of the students' course set history, Avg. Fit Recomm stands for the average fitness of the recommended course sets, Avg. target diff describes the average difference of achieved number of soft skill targets.

Standard deviations are in parentheses. Average satisfied target difference is analysed with paired t-tests: ***$p<0.001$, **$p<0.01$, *$p<0.05$

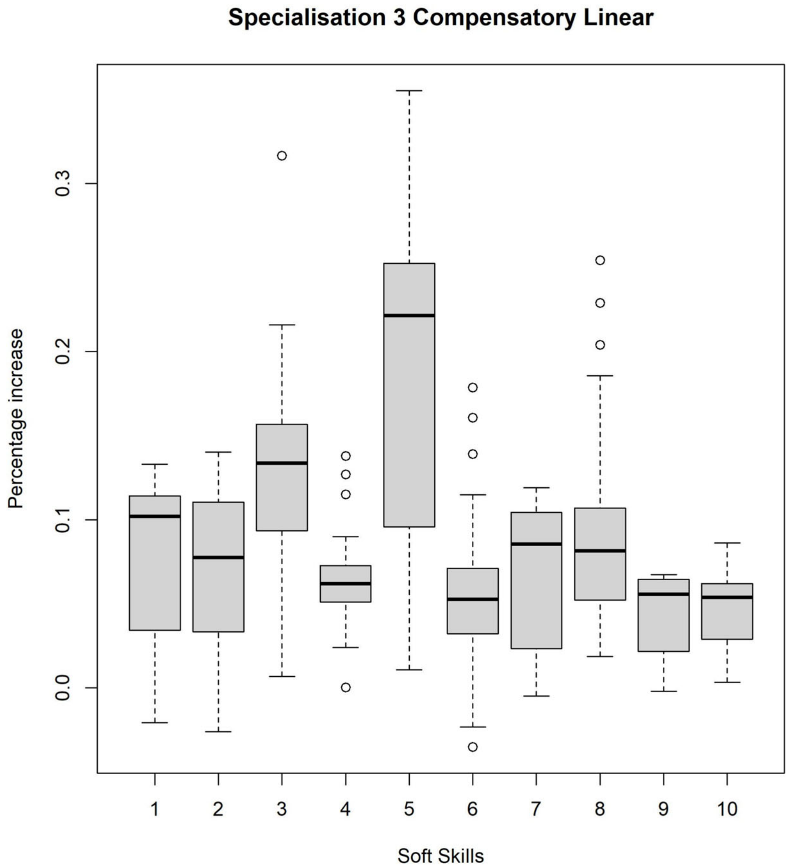

Fig. 9 Soft skill score increase on Specialisation 3 for the Compensatory Linear Fitness. *Note:* The numbers on the horizontal axis corresponds to the soft skills identifiers

Discussion and Limitations

The results suggest that on average the proposed course sets provide in general a higher fitness compared to the real data. This further entails the recommendations may help students better develop their soft skills in order for them to approach to the targeted soft skill profile. The fitness is the hardest part in an optimisation algorithm, especially in this case, where there are 10 assessed soft skills. The different combinations of score functions and aggregation methods represent different pedagogical postures, whether soft skills can be compensated across dimensions, whether bonuses, penalties or none of those should be applied in course sets that allow student to pass the targeted profile. This also implies that we do not intend to conclude which posture is best, but rather remark that comparisons between them may not be feasible since they are based on different stances. For instance, the Logistic score function appears to be the fitness function that delivers the recommendations which most frequently produce an increase, both in the difference of soft skills targets that are passed as well as the average fitness increase change. This is probably due to the strict bonus and penalty policies applied in order to choose the course sets that are best to the students. Nevertheless, this does not mean that this or other functions that convey a similar stance are preferable, since the optimal function depends on the aims of the HEI. It is also important to remark that the considerable differences in the relative increase (e.g., 15% to 390%) in compensatory to partially compensatory aggregations are due to their mathematical definition. For instance, an increase of 10% in each of the 10 scores, would convey a 10% increase in compensatory aggregation fitness (average) whereas it would translate to 159.4% in partially compensatory ($159.4\% = 1.10^{10} - 1$). Therefore, these increases should be compared within the same aggregation method.

In Appendix D, the Figs. 10- 13 depict the boxplots of the soft skill scores for each of the specialisations across the different score functions and aggregation methods. These results show the percentage increase of the same score functions to be similar across aggregation methods, which may suggest that in this case the aggregation method does not greatly influence the aggregated fitness. This is further supported by the fact that the average differences of satisfied targets is relatively independent of whether or not we allow for compensation. This behaviour may not necessarily repeat itself in other cases, but it motivates us to further test this with other datasets.

Our proposed approach could also, in principle, be applied to other types of skills such as hard skills. Nevertheless, hard skills tend not to be as generic as soft skills, and usually do not extend further from the technical expertise of a course. This also entails that most hard skills would only be affected by the courses that intent on developing them. Moreover, the number of hard skills could indeed be considerably larger than their soft skills counterpart. It is important to remark that if our focus was set on hard skills, the interest of a recommender system of course sets would decrease given that there would not be many hard skills that could be developed across various courses. In practice, this would mean that we would need to check which hard skills are more important for the profile, and recommend the courses that develop them.

Limitations, as well as possible extensions, of our approach mostly come from the psychometric model. For instance, our model currently does not account for the effect of specialisations or the effect of the internships, including how the specialisations specifically affect soft skills or the degrees of how different internships promote soft skill development. Ultimately, one could deepen the granularity and extend the current approach to account for other characteristics, such as student interactions within learning tasks in each course. Although, we would need to have these data for each of the courses in order to estimate the effects of these learner interactions.

The current course effects are assumed to be the same across both time and students, even if there are changes regarding lecturers or syllabus, which may not completely realistic. Furthermore, the course sets are assumed not to have sequence-related effects. That is, the effect of having taken c_i before c_j is assumed to be the same as the opposite. While courses are usually taken in parallel in a semester, we could check the sequence by semester. Nevertheless, that would require us to also collect soft skill proficiency measures more than once per year, which can be quite difficult to do for the whole curricula of entire cohorts in HEIs.

As we mentioned previously, there are also limitations regarding the dataset size, which is highly sparse and incompatible for more common RS approaches. This motivated us to use multiple imputation to fill in the missing data from cohort 2021 at stage 0, in order to have sufficient data (1910 data points) for model convergence. We fitted our models to each of the imputed datasets, and pooled the parameter estimates, without the possibility of reserving a subset of the training set for a test set.

Another limitation comes from the students' choices over their courses. This may have an impact on comparability because the students chose their own courses instead of being "placed" in the courses, as in more common experiments where casual inference is studied. In essence, students may have chosen a particular course because they were interested in the specific content of the course, and this interest may affect their performance in the course (by putting more effort compared to another less involved student), creating a confounding effect. In addition, if the system tries to recommend course sets to new students (e.g., students who transfer from other HEIs), the recommendations would be general (per specialisation). This is because the individual random effects for these new students would be equal to zero in the psychometric model. This situation is called the cold-start problem. For the moment our model does not consider the change of student behaviour and performance across time. This is mainly due to the low number of interactions across time (once per year), though with more data, one could adapt the model to account for variability across time.

Aside from the model constraints, there are also computational time limitations, which motivated us to use computer clusters. In this way, the scripts were executed in parallel to reduce the time for the results and the analysis.

Conclusion and Future Work

In this paper we answered **RQ1:** *How can we recommend courses based on the students' soft skills proficiency having short-term small-size datasets from Higher Education?* by proposing a genetic algorithm recommender system framework with fitness functions that represent the aims and stances of HEIs. We showed that a psychometric approach can indeed be implemented as a formal model into genetic algorithms to produce optimised and personalised curricula, driving the exploration of the search space. We also addressed **RQ2:** *How can we represent and leverage both pedagogical postures and soft skills expectations in order to quantify the relevance of a set of course across multiple soft skills?* by evaluating and displaying different score functions and aggregation methods in order to quantify the fitness of a course set across 10 assessed soft skill dimensions. In doing so, we explored and delved a bit further into the course recommendation problem integrating various soft skills as core information towards helping the students in their course choice decision making process. Moreover, this approach, being situated in the common ground between differing research fields, opens a new prospect into the subject of personalised curriculum. We hope this work inspires other researches into investigating other ways to help the students improve their soft skill proficiency across their studies.

There are several ideas for future work. First, one could add course characteristics into the logistic model. For instance, the learning methodologies applied (e.g., problem-based, project-based, inverse classrooms). In that way, we would consider not only the effect of the course as a whole, but also the effects of different methodologies.

Second, the models can be extended in order to account for possible confounding variables related to the interest of the students for specific courses, and its inclusion in the recommendation approach. Third, we could further investigate the targeted soft skill profile, which may be built amongst various stakeholders (e.g., the academic institution, the industry, alumni and even current students) in order for each partner to contribute to the overall expected reference from the engineering students at the end of their studies. This could also include the definition and fine-tuning of weights of each contributor, as well as the definition of intervals of acceptance instead instead of fixed points. Fourth, course sequencing (i.e., the effect of following a particular course after another) could be considered, though that would also imply modifying the model that predicts soft skill proficiency. Fifth, we could test whether the aggregation method changes considerably the results in terms of the final recommended course sets. Sixth, we could explore this problem considering a multi-objective perspective, where we would try to maximise each soft skill score, instead of the overall fitness of the course set. Lastly, an user profile-inspired score function, based on historical data and a more data-driven approach, could be used instead of functional forms as in this paper.

Appendix A: Appendix

A.1 Soft Skill Assessment Scheme

Table 8 Soft skill Assessment Scheme

Skill	Description	Mark
Problem Solving	Solve problems in familiar environments or in known contexts. Use solutions that are already outlined	1
	Solve related problems in new environments or in unknown contexts. Find solutions adapted to each situation	2
	Solve problems that are not always well defined, in complex environments or in contexts subject to strong constraints	3
	Reformulate problems according to different constraints and find adapted and efficient solutions	4
Innovation and Creativity	Contribute to the search for ideas and solutions by participating in exchanges and creativity sessions and propose improvements	1
	Propose and apply proven solutions to new or different contexts	2
	Design and implement new solutions with a view to efficiency	3
	Recognise, transmit and implement the conditions and processes for generating innovation	4
Organisation	Work independently on the basis of indications and instructions given. Monitor performance indicators, detect and report problems within the scope of activity	1

Table 8 continued

Skill	Description	Mark
	Prioritise and plan their own workload, evaluate and correct completed activities (use performance indicators to assist in decision making)	2
	Prioritise and establish actions according to the stakes of the activities. Set up new relevant indicators. Share and promote best practices	3
	Transmit and share methods of organisation and rigour with his interlocutors. Encourage them to use and follow relevant performance indicators. Deploy continuous improvement plans	4
Agility and adaptability	Implement activity changes that are requested	1
	Adapt and re-prioritise their activities and organisation to changes and constraints	2
	Evaluate the impact of changes and propose appropriate responses or solutions	3
	Anticipate future developments and changes	4
Interaction and Communication	Listen actively, express and formalise a point of view clearly, share information. Reformulate an idea without distorting it	1
	Effectively present an argument in a logical and argued manner, both in writing and orally. Know how to use a vocabulary that goes beyond the usual	2
	Use expression techniques (written and oral) adapted to the message to be delivered and the target audience (specialists and/or non-specialists), in a clear and unambiguous manner	3
	Communicate skilfully (vocabulary, style, ...) and finely in complex situations (sensitive message, difficult audience, unexpected situation...)	4
Project Management	Work within a group (a team) and collaborate with team members in an open manner by communicating feedback on work	1
	Manage a small independent project or a part within a larger program. Accompany one or two collaborators on an operational activity	2
	Lead a major project or coordinate several operational projects simultaneously. Lead a complete team on an operational activity or project	3
	Lead and coordinate several strategic or operational projects simultaneously, training teams on project management, setting up adapted animation devices. Settling conflicts and arbitration situations in an objective and factual way without breaking interpersonal relations	4
Conviction	Explain a point of view in a clear way and with predefined or prepared arguments. Listen to, understand and reproduce a need expressed by others	1

Table 8 continued

Skill	Description	Mark
	Adapt one's behaviour, attitude, speech and arguments according to the audience in order to maximise the quality of exchanges. Interact to reformulate and deepen one's need in order to specify it and propose a response	2
	Identify and decipher the positions of the different strategic audiences, anticipate their expectations and reactions, identify and reach the right influence relays with the people to be convinced. To be a force of proposal in relation to an expressed need while rallying the stakeholders	3
	Implement strategic actions in a complex environment in a relevant and recurring manner to convince and influence key players. Anticipate the needs of their "lients" their environments and guide them in their evolution	4
Stress Management	Work in low stress situations	1
	Adapt themselves to temporary stressful situations	2
	Adapt to prolonged stress	3
	Use strategies to deal with stress themselves and for their collaborators	4
Leadership and influence	To take a step back and take initiatives in the service of activities and collaborators belonging to a close circle	1
	Share their own vision with familiar and occasional (or close and temporary) collaborators	2
	Promote their vision to internal and external decision-makers, and encourage their teams to take the initiative	3
	Pool resources and partners by creating a dynamic around a strategy and/or a change process	4
Decision Making	Makes decisions based only on rules	1
	Makes decisions in situations where rule interpretation is possible	2
	Makes decisions interpreting the rules and improves them	3
	Makes complex decisions in absence of rules	4

From Pinos Ullauri et al. (2024)

Appendix B: Score Functions

B.1 Notation

- $minsskill$ is the minimum skill soft skill proficiency level, $minsskill = 1$.
- $maxsskill$ is the maximum skill soft skill proficiency level, $maxsskill = 4$.
- $\overline{sskill_{is}} \in [1, 4]$ is the expected soft skill proficiency of soft skill i from student s at the end of the program (stage 2).
- $tarProf_{im} \in [1, 4]$ is the targeted proficiency at soft skill i from students of the specialisation m.
- sf_{imin} is the minimum value of function f given soft skill i and $minsskill$.
- sf_{imax} is the maximum value of function f given soft skill i and $maxsskill$

- $sf_{crossing}$ is a mathematical expression which allows all functions to cross the point where $\overline{sskill}_{is} = tarProf_{im}$ in the linear function.

B.2 Linear Score Function

The linear function, as with the following score functions, are scaled in the y-axis. Consequently, the scores are continuous values in the range between zero and one. Moreover, since the constraints are fixed, the slope remains the same, regardless of the targeted proficiency $tarProf_{im}$.

$$Linear_{is} = \frac{\overline{sskill}_{is} - minsskill}{maxsskill - minsskill} \tag{B1}$$

B.3 Logistic Sigmoid Score Function

The logistic sigmoid function normally returns values within the range of (-1,1). Since the scores are required to be in the range [0,1], we leverage a mathematical term $sf_{crossing}$ in order to displace and re-scale the function. This displacement affects the function both vertically and horizontally. It makes it cross at the point where $\overline{sskill}_{is} = tarProf_{im}$ in the linear function. Moreover, the function is re-scaled, which makes the scores to be indeed within the range between zero and one. The exponent sets the inclination for the logistic function. The higher the exponent, the more the function resembles a Heaviside step function. The lower the exponent, the more the function flattens becoming at some point a horizontal line. We set the exponent at 3, to portray the stricter scoring compared to the linear function.

$$sf_{crossing} = \frac{tarProf_{im} - minsskill}{maxsskill - minsskill} \tag{B2}$$

$$Logistic_{is} = \frac{1}{1 + (\frac{1 - f_{crossing}}{f_{crossing}})e^{3(tarProf_{im} - \overline{sskill}_{is})}} \tag{B3}$$

B.4 Quadratic Score Function

The quadratic function is a piece-wise function, composed of a negative quadratic part on the left side ($\overline{sskill}_{is} < expProf_{im}$), and a positive quadratic function when the targeted threshold is passed. Similarly to the sigmoid function, a quadratic function leverages the mathematical term $sf_{crossing}$ to displace the function. Nonetheless, it also uses sf_{min} and sf_{max} for the scaling. This makes the function returns scores in the range [0,1].

$$\begin{aligned} sf_{crossing} &= \frac{expProf_{im} - minsskill}{maxsskill - minsskill} \\ sf_{min} &= -(minsskill - expProf_{im})^2 \\ sf_{max} &= (maxsskill - expProf_{im})^2 \end{aligned} \tag{B4}$$

$$Quadratic_{is} = \begin{cases} \frac{f_{crossing}\overline{sskill}_{is} - expProf_{im}{}^{2}}{f_{min}} + f_{crossing} & \text{if } \overline{sskill}_{is} < expProf_{im} \\ \frac{(1-f_{crossing})\overline{sskill}_{is} - expProf_{im}{}^{2}}{f_{max}} + f_{crossing} & \text{if } \overline{sskill}_{is} \geq expProf_{im} \end{cases} \quad \text{(B5)}$$

Appendix C: Genetic Algorithm

C.1 Pseudo-code

The uniform crossover essentially copies one of the parents into the child course set. Then, each course is tested with a 50% probability to see if it continues or is replaced by the course of the other parent. This process also verifies if a course is already in the set to avoid duplicates.

Algorithm 1 Uniform Crossover.

Require: Parent course sets p_1, p_2. Crossover rate x
Ensure: Adequate child creation from parents p_1 and p_2

```
child ← default copy from parent course set p2
r_x ← randomly generated number between [0, 1]
if r_x ≤ x then
    for i in p1 do
        r ← randomly generated number between [0, 1]
        if p1 at index i is not within p2 then
            if r > 0.50 then
                child at index i ← copy p1 at index i
            end if
        end if
    end for
end if
return child
```

The mutation process generates a new course in the course set to add randomness in the process. Likewise to the crossover algorithm, it validates and avoids duplicates in the course sets.

C.2 Source Code

Please refer to https://github.com/F-FIDELO-19-008-FLEURY/course-recommender.git, for the implementation of the genetic algorithm used in this study. Furthermore, the source code has also been archived in Software Heritage, and can be accessed via the link: https://archive.softwareheritage.org/browse/snapshot/bf7fd2bc2d3f35378e04541f69e2c95ba967e971/directory/?origin_url=https://github.com/F-FIDELO-19-008-FLEURY/course-recommender.

Algorithm 2 Mutation.

Require: Ensemble g of course sets $p_1, p_2, \ldots, p_{100}$. Mutation rate m
Require: List of possible courses l
Ensure: Appropriate mutated course set
  $r_m \leftarrow$ randomly generated number between [0, 1]
  if $r_m \leq m$ **then**
    for p_i in g **do**
      available courses $\leftarrow$ courses from l not in p_i
      if c_1 at index i is not within c_2 **then**
        if $r > 0.50$ **then**
          $course\ child\ at\ index\ i \leftarrow\ copy\ c_1\ at\ index\ i$
        end if
      end if
    end for
  end if
  return $child$

Appendix D: Percentage Increase Across Soft Skills

It can be seen that soft skill 5 is quite favoured in both compensatory and partially compensatory methods. Moreover, the Quadratic fitness function highly profits from soft skill 5 to compensate the lack of high increases on the other soft skills.

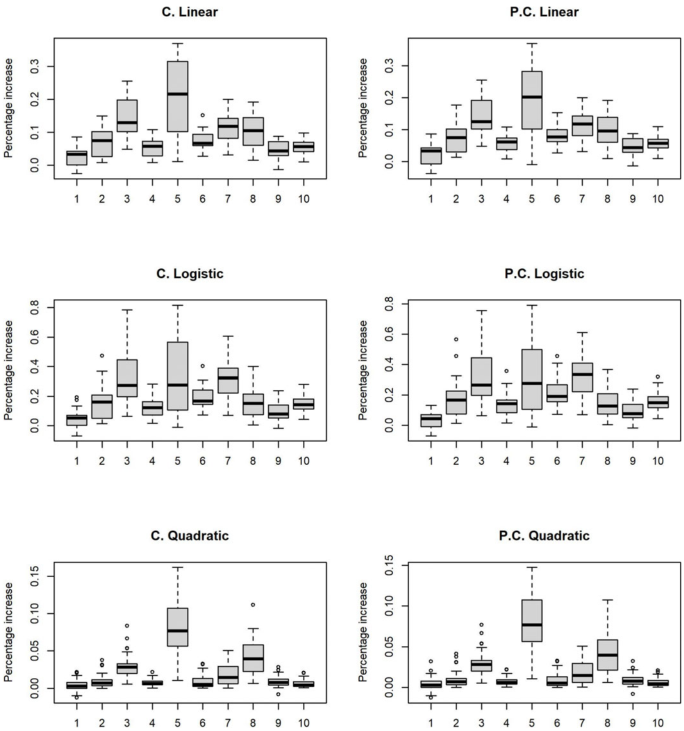

Fig. 10 Soft skill score increase on Specialisation 1. Note: The C. stands for Compensatory whereas P. C. for Partially Compensatory. Also, the numbers on the horizontal axis corresponds to the soft skills identifiers

Figure 11 shows the percentage increases on Specialisation 2, which compared to Specialisation 1, show a different behaviour of the scores. For the Logistic function, there are not prominent favoured soft skills, whereas for the Quadratic fitness soft skill 8 is highly increased.

Specialisation 3 shows a similar variability as with Specialisation 1, with emphasis on some soft skills such as soft skill 1, 5 and 8. The Quadratic fitness shows a considerable difference of the increases between the soft skills.

Finally, Specialisation 4 shows a more balanced increase across soft skills compared to the other specialisations. This, however, does not occur for the Quadratic functions, which rely on soft skill 5 and 8 to compensate partially or not the other soft skills.

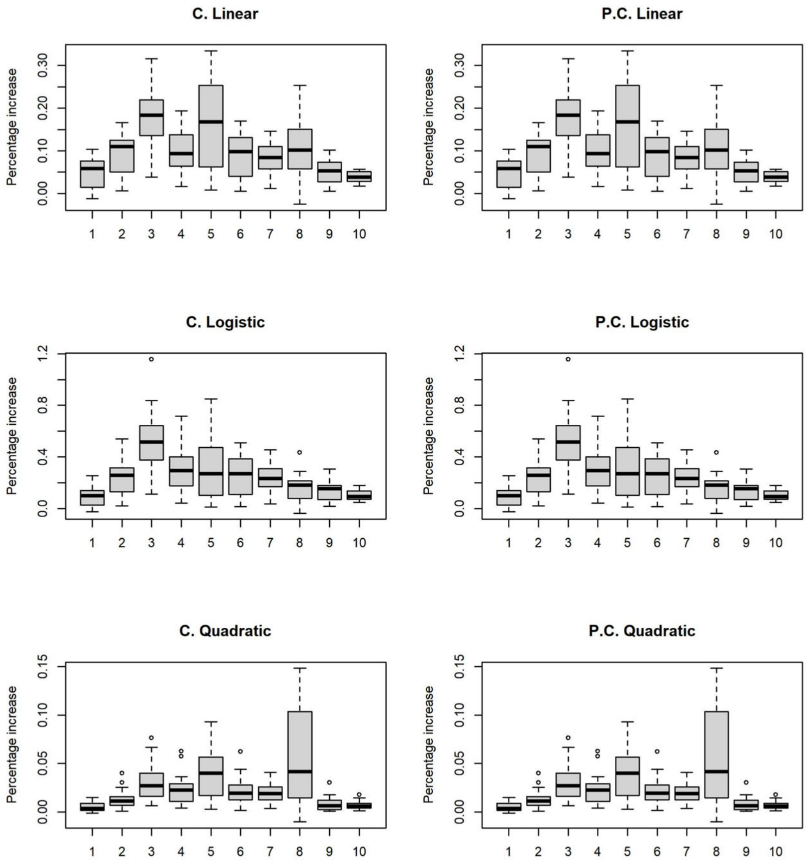

Fig. 11 Soft skill score increase on Specialisation 2. Note: The C. stands for Compensatory whereas P. C. for Partially Compensatory. Also, the numbers on the horizontal axis corresponds to the soft skills identifiers

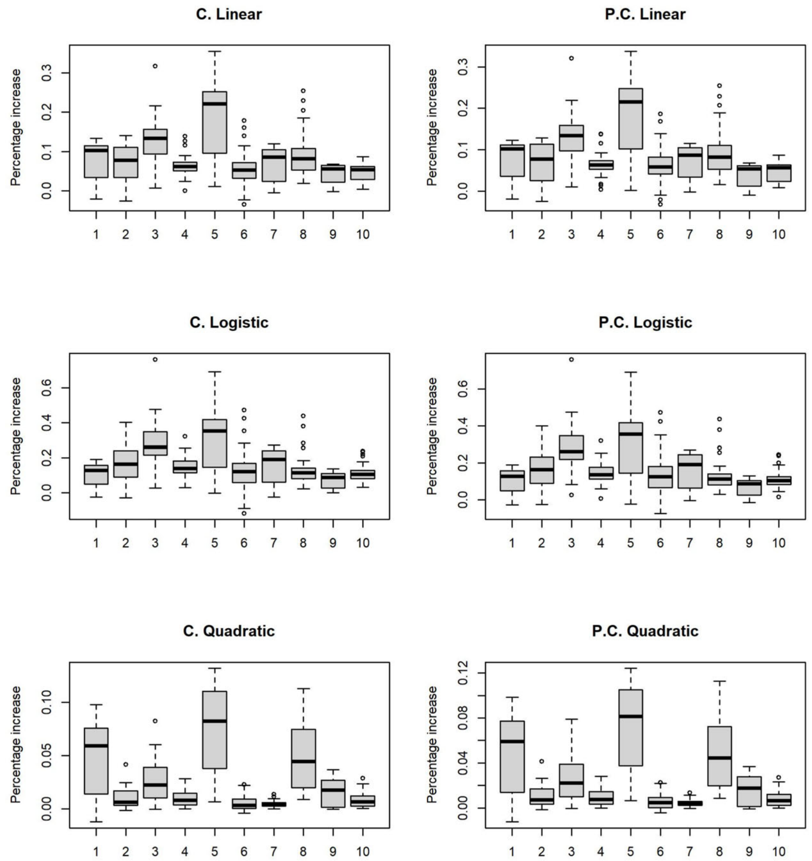

Fig. 12 Soft skill score increase on Specialisation 3. Note: The C. stands for Compensatory whereas P. C. for Partially Compensatory. Also, the numbers on the horizontal axis corresponds to the soft skills identifiers

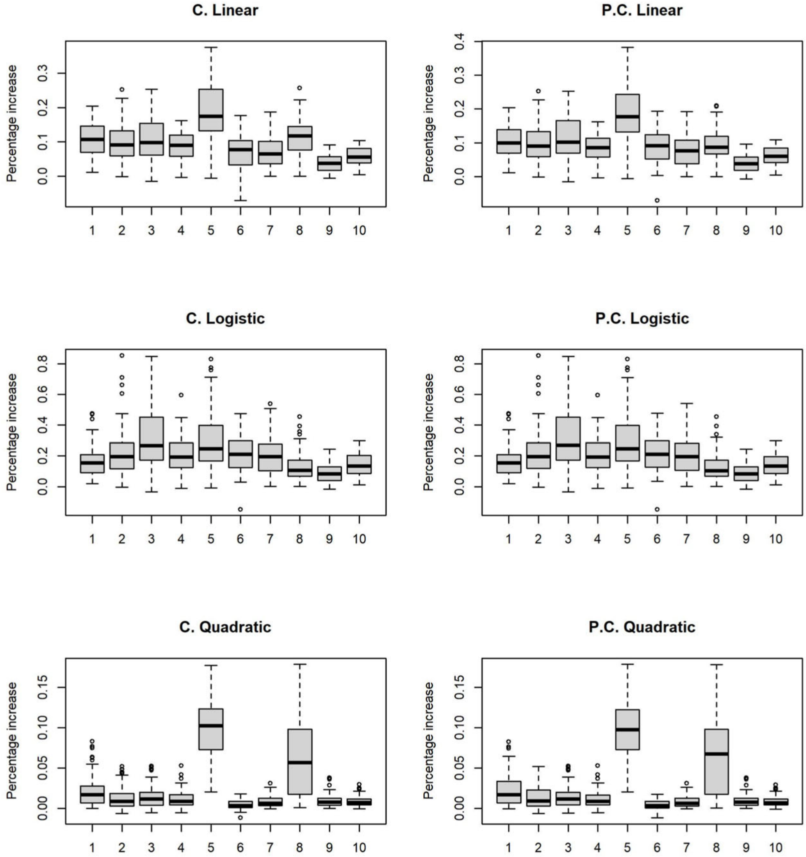

Fig. 13 Soft skill score increase on Specialisation 4. Note: The C. stands for Compensatory whereas P. C. for Partially Compensatory. Also, the numbers on the horizontal axis corresponds to the soft skills identifiers

Acknowledgements We wanted to thank Jerémie Humeau for his help with the computer clusters used for the running of the genetic algorithms. This work was funded by I-SITE ULNE (42 Rue Paul Duez, 59800 Lille, France) under the project SUCCESS. The SUCCESS project aims on supporting the assessment of soft skill proficiency on postgraduate students, and provide personalised recommendations of courses based on their soft skills.

Author Contributions L. A. P. U. is the main author of this article, contributing in all stages: conceptualisation, methodology, implementation of the algorithms, validation, formal analysis, visualisations, and writing of the original and consequent drafts. A. L. contributed to the conceptualisation, validation, and review of the manuscript. A. K. contributed to the validation and review of the manuscript. M. V. contributed to the validation and review of the manuscript. A. F. contributed to the funding acquisition, validation, and review of the manuscript. W. V. D. N. contributed to the methodology, formal analysis, and review of the manuscript. All authors read and approved the final manuscript.

Funding This work was funded by i-SITE ULNE (42 Rue Paul Duez, 59800 Lille, France) under the project SUCCESS and grant F-FIDELO-19-008-FLEURY. The SUCCESS project aims on supporting the assessment of soft skill proficiency on postgraduate students, and provide personalised recommendations of courses based on their soft skills.

Data Availability The authors agree to share the data used to train and test the algorithms in this paper upon reasonable request.

Declarations

Competing interests The authors declare no competing interests.

References

Adams, R. J., Wilson, M., & Wang, W.-C. (1997). The multidimensional random coefficients multinomial logit model. *Applied Psychological Measurement, 21*(1), 1–23. https://doi.org/10.1177/0146621697211001

Alhijawi, B., & Kilani, Y. (2020). A collaborative filtering recommender system using genetic algorithm. *Information Processing & Management, 57*(6), 102310. https://doi.org/10.1016/j.ipm.2020.102310

Almonte, R. (2021). A practical guide to soft skills: Communication, psychology, and ethics for your professional life. Routledge, Taylor & Francis Group, ???. https://doi.org/10.4324/9781003212942

Al-Muhaideb, S., & Menai, M. E. B. (2011). Evolutionary computation approaches to the curriculum sequencing problem. *Natural Computing, 10*, 891–920. https://doi.org/10.1007/s11047-010-9246-5

Arvanitis, A., Touloumakos, A. K., Dimitropoulou, P., Vlemincx, E., Theodorou, M., & Panayiotou, G. (2022). Learning how to learn in a real-life context: Insights from expert focus groups on narrowing the soft-skills gap. *European Journal of Psychology Open, 81*, 71–77. https://doi.org/10.1024/2673-8627/a000027

Barrie, S. (2006). Understanding what we mean by generic attributes of graduates. *Higher Education, 51*, 215–241. https://doi.org/10.1007/s10734-004-6384-7

Barrie, S. C. (2007). A conceptual framework for the teaching and learning of generic graduate attributes. *Studies in Higher Education, 32*(4), 439–458. https://doi.org/10.1080/03075070701476100

Bobadilla, J., Ortega, F., Hernando, A., & Alcalá, J. (2011). Improving collaborative filtering recommender system results and performance using genetic algorithms. *Knowledge-Based Systems, 24*(8), 1310–1316. https://doi.org/10.1016/j.knosys.2011.06.005

Cacciolatti, L., Lee, S. H., & Molinero, C. M. (2017). Clashing institutional interests in skills between government and industry: An analysis of demand for technical and soft skills of graduates in the UK. *Technological Forecasting and Social Change, 119*, 139–153. https://doi.org/10.1016/j.techfore.2017.03.024

Chamorro-Premuzic, T., Arteche, A., Bremner, A., Greven, C., & Furnham, A. (2010). Soft skills in higher education: Importance and improvement ratings as a function of individual differences and academic performance. *Educational Psychology, 30*, 221–241. https://doi.org/10.1080/01443410903560278

Deist, F. D. L., & Winterton, J. (2005). What is competence? *Human Resource Development International, 8*(1), 27–46. https://doi.org/10.1080/1367886042000338227

Dwivedi, P., Kant, V., & Bharadwaj, K. K. (2018). Learning path recommendation based on modified variable length genetic algorithm. *Education and Information Technologies, 23*, 819–836. https://doi.org/10.1007/s10639-017-9637-7

European Higher Education Area (2016). Bologna process: Employability of graduates. https://www.ehea.info/cid102525/employability-of-graduates.html

Farndale, E., Scullion, H., & Sparrow, P. (2010). The role of the corporate hr function in global talent management. *Journal of World Business, 45*(2), 161–168. https://doi.org/10.1016/j.jwb.2009.09.012. Global Talent Management

Gad, A. F. (2023). Pygad: An intuitive genetic algorithm python library. *Multimedia Tools and Applications*. https://doi.org/10.1007/s11042-023-17167-y

Guruge, D. B., Kadel, R., & Halder, S. J. (2021). The state of the art in methodologies of course recommender systems—a review of recent research. *Data, 6*(2). https://doi.org/10.3390/data6020018

Hallinan, B., & Striphas, T. (2016). Recommended for you: The netflix prize and the production of algorithmic culture. *New Media & Society, 18*(1), 117–137. https://doi.org/10.1177/1461444814538646

Han, J.-w, Jo, J.-c, Ji, H.-s, & Lim, H.-s. (2016). A collaborative recommender system for learning courses considering the relevance of a learner's learning skills. *Cluster Computing, 19*, 2273–2284. https://doi.org/10.1007/s10586-016-0670-x

Hanson, A., Bullers, K., Howard, A., Polo, R., Tomlinson, S., & Orriola, J. (2019). Using a reflexive process to investigate organizational change: The use of the research spider matrix. *Medical Reference Services Quarterly, 38*, 312–325. https://doi.org/10.1080/02763869.2019.1657724

Hu, L. (2024). Programming and 21st century skill development in k-12 schools: A multidimensional meta-analysis. *Journal of Computer Assisted Learning, 40*(2), 610–636. https://doi.org/10.1111/jcal.12904

Hurrell, S. A. (2016). Rethinking the soft skills deficit blame game: Employers, skills withdrawal and the reporting of soft skills gaps. *Human Relations, 69*(3), 605–628. https://doi.org/10.1177/0018726715591636

Hurrell, S. A., Scholarios, D., & Thompson, P. (2013). More than a 'humpty dumpty' term: Strengthening the conceptualization of soft skills. *Economic and Industrial Democracy, 34*(1), 161–182. https://doi.org/10.1177/0143831X12444934

Ibrahim, M. E., Yang, Y., Ndzi, D. L., Yang, G., & Al-Maliki, M. (2019). Ontology-based personalized course recommendation framework. *IEEE Access, 7*, 5180–5199. https://doi.org/10.1109/ACCESS.2018.2889635

José-García, A., Sneyd, A., Melro, A., Ollagnier, A., Tarling, G., Zhang, H., Stevenson, M., Everson, R., & Arthur, R. (2023). C3-ioc: A career guidance system for assessing student skills using machine learning and network visualisation. *International Journal of Artificial Intelligence in Education, 33*(4), 1092–1119. https://doi.org/10.1007/s40593-022-00317-y

Lebis, A., Humeau, J., Fleury, A., Lucas, F., & Vermeulen, M. (2023). Fully individualized curriculum with decaying knowledge, a new hard problem: Investigation and recommendations. *International Journal of Artificial Intelligence in Education*. https://doi.org/10.1007/s40593-023-00376-9

Li, Q., & Kim, J. (2021). A deep learning-based course recommender system for sustainable development in education. *Applied Sciences, 11*(19). https://doi.org/10.3390/app11198993

Llorens Garcia, A., Prat Farran, J. d., & Berbegal-Mirabent, J. (2019). Ict skills gap in spain: Before and after a decade of harmonizing the european higher education area. *Computer Applications in Engineering Education, 27*(4), 934–942. https://doi.org/10.1002/cae.22132

Marcos, L., Martínez, J.-J., & Gutierrez, J.-A. (2008). Swarm intelligence in e-learning: A learning object sequencing agent based on competencies. In *Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation. GECCO '08* (pp. 17–24). Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/1389095.1389099

Muukkonen, H., Lakkala, M., Ilomäki, L., & Toom, A. (2022). Juxtaposing generic skills development in collaborative knowledge work competences and related pedagogical practices in higher education. *Frontiers in Education, 7*. https://doi.org/10.3389/feduc.2022.886726

Nabizadeh, A. H., Gonçalves, D., Gama, S., Jorge, J., & Rafsanjani, H. N. (2020). Adaptive learning path recommender approach using auxiliary learning objects. *Computers & Education, 147*, 103777. https://doi.org/10.1016/j.compedu.2019.103777

Notari, M., Baumgartner, A., & Herzog, W. (2014). Social skills as predictors of communication, performance and quality of collaboration in project-based learning. *Journal of Computer Assisted Learning, 30*(2), 132–147. https://doi.org/10.1111/jcal.12026

Novais, A. S. d., Matelli, J. A., & Silva, M. B. (2023). Fuzzy soft skills assessment through active learning sessions. *International Journal of Artificial Intelligence in Education*. https://doi.org/10.1007/s40593-023-00332-7

Peduzzi, P., Concato, J., Kemper, E., Holford, T. R., & Feinstein, A. R. (1996). A simulation study of the number of events per variable in logistic regression analysis. *Journal of Clinical Epidemiology, 49*(12), 1373–1379.

Pinos Ullauri, L. A., Lebis, A., Karami, A., Vermeulen, M., Fleury, A., & Van Den Noortgate, W. (2024). Modeling the effect of postgraduate courses on soft skills: A practical approach. *Frontiers in Psychology, 14*. https://doi.org/10.3389/fpsyg.2023.1281465

Pinos Ullauri, L. A., Lebis, A., Karami, A., Vermeulen, M., Fleury, A., & Van Ven Noortgate, W. (2023). Système de recommandation de cours basé sur les soft skills: Une approche utilisant les algorithmes génétiques. In *EIAH2023 : 11ème Conférence sur les Environnements Informatiques Pour l'Apprentissage Humain, Brest, France* (pp. 344–347).

Reckase, M. D. (2009). *Multidimensional item response theory* (1st ed). Springer, ???.

Ricci, F., Rokach, L., & Shapira, B. (2010). Introduction to Recommender Systems Handbook. *Recommender Systems Handbook* (pp. 1–35). Springer, ???. https://doi.org/10.1007/978-0-387-85820-3_1

Richa, S., Paul, J., & Tewari, V. (2021). The soft skills gap: A bottleneck in the talent supply in emerging economies. *The International Journal of Human Resource Management, 33*, 1–32. https://doi.org/10.1080/09585192.2020.1871399

Riley, R. D., Snell, K. I., Ensor, J., Burke, D. L., Harrell, F. E., Jr., Moons, K. G., & Collins, G. S. (2019). Minimum sample size for developing a multivariable prediction model: Part ii - binary and time-to-event outcomes. *Statistics in Medicine, 38*(7), 1276–1296. https://doi.org/10.1002/sim.7992

Sancho-Thomas, P., Fuentes-Fernández, R., & Fernández-Manjón, B. (2009). Learning teamwork skills in university programming courses. *Computers & Education, 53*(2), 517–531. https://doi.org/10.1016/j.compedu.2009.03.010

Schultz, C. (2010). Hr competencies at a merged higher education institution. *South African Journal of Human Resource Management, 8*. https://doi.org/10.4102/sajhrm.v8i1.225

Shakhsi-Niaei, M., & Abuei-Mehrizi, H. (2020). An optimization-based decision support system for students' personalized long-term course planning. *Computer Applications in Engineering Education, 28*(5), 1247–1264. https://doi.org/10.1002/cae.22299

Smeden, M., Moons, K. G., Groot, J. A., Collins, G. S., Altman, D. G., Eijkemans, M. J., & Reitsma, J. B. (2019). Sample size for binary logistic prediction models: Beyond events per variable criteria. *Statistical Methods in Medical Research, 28*(8), 2455–2474. https://doi.org/10.1177/0962280218784726. PMID: 29966490.

Son, N. T., Jaafar, J., Aziz, I. A., & Anh, B. N. (2021). Meta-heuristic algorithms for learning path recommender at mooc. *IEEE Access, 9*, 59093–59107. https://doi.org/10.1109/ACCESS.2021.3072222

Succi, C., & Canovi, M. (2019). Soft skills to enhance graduate employability: Comparing students and employers' perceptions. *Studies in Higher Education, 45*, 1–14. https://doi.org/10.1080/03075079.2019.1585420

Tadjer, H., Lafifi, Y., Derindere, M., Gülseçen, S., & Seridi-Bouchelaghem, H. (2018). What are the important social skills of students in higher education?

Tadjer, H., Lafifi, Y., Seridi-Bouchelaghem, H., & Gülseçen, S. (2022). Improving soft skills based on students' traces in problem-based learning environments. *Interactive Learning Environments, 30*(10), 1879–1896. https://doi.org/10.1080/10494820.2020.1753215

Tesch, D., Jiang, J. J., & Klein, G. (2003). The impact of information system personnel skill discrepancies on stakeholder satisfaction. *Decision Sciences, 34*(1), 107–129. https://doi.org/10.1111/1540-5915.02371

Tuononen, T., Hyytinen, H., Kleemola, K., Hailikari, T., Männikkö, I., & Toom, A. (2022). Systematic review of learning generic skills in higher education-enhancing and impeding factors. *Frontiers in Education, 7*. https://doi.org/10.3389/feduc.2022.885917

Verbeke, G., Fieuws, S., Molenberghs, G., & Davidian, M. (2014). The analysis of multivariate longitudinal data: A review. *Statistical Methods in Medical Research, 23*(1), 42–59. https://doi.org/10.1177/0962280212445834

Yacobson, E., Toda, A. M., Cristea, A. I., & Alexandron, G. (2024). Recommender systems for teachers: The relation between social ties and the effectiveness of socially-based features. *Computers & Education, 210*, 104960. https://doi.org/10.1016/j.compedu.2023.104960

Authors and Affiliations

Luis Alberto Pinos Ullauri[1,2,3] · Alexis Lebis[3] · Abir Karami[4] · Mathieu Vermeulen[3] · Anthony Fleury[3] · Wim Van den Noortgate[1,2]

✉ Luis Alberto Pinos Ullauri
luis.pinos.research@proton.me

Alexis Lebis
alexis.lebis@imt-nord-europe.fr

Abir Karami
abir.karami@univ-catholille.fr

Mathieu Vermeulen
mathieu.vermeulen@imt-nord-europe.fr

Anthony Fleury
anthony.fleury@imt-nord-europe.fr

Wim Van den Noortgate
wim.vandennoortgate@kuleuven.be

1 Faculty of Psychology and Educational Sciences, KU Leuven, 51-53 Etienne Sabbelaan, Kortrijk 8500, Belgium

2 imec research group ITEC, KU Leuven, 51-53 Etienne Sabbelaan, Kortrijk 8500, Belgium

3 Centre for Digital Systems, IMT Nord Europe, 764 Boulevard Lahure, Douai 59500, France

4 ICL Junia, Université Catholique de Lille, LITL, Lille F-59000, France