



Impressions and Strategies of Academic Advisors When Using a Grade Prediction Tool During Term Planning

Gonzalo Gabriel Méndez, Luis Galárraga, Katherine Chiluiza, Patricio Mendoza

► To cite this version:

Gonzalo Gabriel Méndez, Luis Galárraga, Katherine Chiluiza, Patricio Mendoza. Impressions and Strategies of Academic Advisors When Using a Grade Prediction Tool During Term Planning. CHI 2023 - Conference on Human Factors in Computing Systems, Apr 2023, Hamburg, Germany. pp.1-18, 10.1145/3544548.3581575 . hal-04132566

HAL Id: hal-04132566

<https://inria.hal.science/hal-04132566>

Submitted on 26 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Impressions and Strategies of Academic Advisors When Using a Grade Prediction Tool During Term Planning

Gonzalo Gabriel Méndez

gmendez@espol.edu.ec

Escuela Superior Politécnica del Litoral
Guayaquil, Ecuador

Katherine Chiluiza

kchilui@espol.edu.ec

Escuela Superior Politécnica del Litoral
Guayaquil, Ecuador

Luis Galárraga

luis.galarraga@inria.fr

INRIA

Rennes, France

Patricio Mendoza

pdmendoz@espol.edu.ec

Escuela Superior Politécnica del Litoral
Guayaquil, Ecuador

ABSTRACT

Academic advising brings numerous benefits to the mission of Higher Education Institutions. One central and challenging duty of advisors is course recommendation for term planning. This task requires both knowledge of the study programs as well as a thorough analysis of the students' unique circumstances. Limited time and a large student population make this task overwhelming. As a result, an important body of research has sought to expedite term planning via data-oriented decision-support tools. The impact of such tools on students has been extensively studied. However, the advisors' perspective remains largely unexplored. We contribute to redressing this gap by studying how a grade prediction tool shapes academic advisors' approach to course recommendation. We found that while the advisors' usual strategies tend to prevail, their recommendations largely depend on the advisee's historical performance. That said, advisors also acknowledge the limitations of grades as a measure of academic success.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in visualization**.

KEYWORDS

academic advising, course recommendation, grade prediction

ACM Reference Format:

Gonzalo Gabriel Méndez, Luis Galárraga, Katherine Chiluiza, and Patricio Mendoza. 2023. Impressions and Strategies of Academic Advisors When Using a Grade Prediction Tool During Term Planning. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3544548.3581575>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9421-5/23/04...\$15.00

<https://doi.org/10.1145/3544548.3581575>

1 INTRODUCTION

Academic advising processes are implemented by Higher Education Institutions (HEIs) to promote exchanges between students and experts in aspects of the students' academic life [27]. These exchanges are conducted under the assumption that the experts' advice contributes to the creation of supportive academic environments that meet the students' needs and, hence, foster their development. A vast body of research on academic support systems has identified multiple benefits of academic advising both for students and HEIs. For students, advising guarantees not only more informed enrollment decisions, but also guidance on matters that may affect their academic success and performance [6, 20]. For educational institutions, academic advising can provide first-hand feedback on the impact of institutional policies, keep students focused on relevant goals, channel the students' academic interests, and establish policies to improve student achievement [14, 19]. For all these reasons, academic advising is often a requirement of educational accreditation boards [5] such as ABET¹.

At the center of academic advising lie the academic advisors, whose role is fundamental for the success of any academic support system. Advisors act as mediators that inform students about institutional policies, principles, and expectations. By meeting students on a regular basis, advisors can identify institutional features that adversely affect the students' experience. Equally important, advisors get to develop a sense of the students' needs beyond academic matters. All this can promote the development of institutional actions oriented to improve student well-being.

Although widely effective, some aspects of academic advising often require significant logistical efforts and resources. In particular, guiding students in the selection of courses for a forthcoming academic term can be prone to errors—due to the introduction of subjective judgments. This task requires assessing each student's academic history in light of their enrollment options at a specific point of their study program, and is particularly challenging in programs with curricula that do not include many elective courses. The challenges around course recommendation have motivated several research efforts to relieve the advisors' workload. Most of these efforts have focused on empowering students with data and tools to make faster and more informed enrollment decisions. The effects of these tools on the students' decision-making process and

¹<https://www.abet.org>

final decisions have been extensively investigated and are quite well understood.

On the contrary, advisor-oriented support is less common, and the effects of data-based tools on the advisors' approach to course recommendation are lesser-known. In this paper we take steps to fill this gap by investigating how a tool that provides course grade predictions (and explanations thereof) shapes the approach of academic advisors to course recommendation. We present the findings of a study in which academic advisors of an Engineering-oriented public university used a grade prediction tool to recommend a set of courses to students with three different academic profiles, namely low-, average-, and high-performing. We analyzed the effects of the tool on the advisors' decisions as well as on their recommendation strategy. Moreover, we studied how these decisions differ based on the advisee's performance profile. Our observations suggest that when using the tool, advisors barely change their self-reported usual recommendation approach. That is, the grade predictions and explanations provided by the tool did not have a major influence on how advisors decided which courses students should take in an upcoming term. That said, advisors did look at complementary information such as the student's academic and enrollment history, in order to build a more precise profile of the advisee. This, in particular, resulted in longer interactions with the tool when making recommendations for the low-performing student, for whom advisors also tried to maximize the GPA gain. We also found that a grade prediction tool may rise concerns not only about the trustworthiness of the predictions and their impact on the students' motivations, but also about the limitations of the GPA to characterize academic performance.

With our study results, we contribute: 1) a characterization of the advisor's decision-making process for course recommendation in the presence of grade predictions and different student academic profiles; 2) a discussion of the advisors' concerns raised by a grade-prediction tool in the context of course recommendation; and 3) a reflection on the need for human intervention to prevent the undesired effects of grade-based predictions in academic support systems. Our findings provide guidance for the design of data-based solutions to make the student-advisor dialogue more effective.

2 RELATED WORK

Before describing our study, we first summarize the challenges around academic advising that motivate the use of data-oriented decision-support tools. We then survey studies on the effects of data and predictive tools on students and advisors. We conclude this section with a body of empirical evidence that highlights several features of AI-based tools that impact people's decision making, particularly in the context of academic advising and course recommendation. This body of work constitutes the conceptual base that informs our study.

2.1 Academic Advising and the Challenges of Course Recommendation

Appropriate academic advising is critical for the success of students [6, 20, 61]. Advising helps HEIs in supporting at-risk students and strengthening their motivations [29] which, in turn, can reduce student attrition and prevent graduation delays and potential

dropouts [57, 60]. The Global Community for Academic Advising (NACADA²) recognizes four dimensions of academic advising at HEIs: (i) *who* does the advising (e.g., faculty vs. other dedicated professionals), (ii) *where* and *how* advising services are available (e.g., on-campus vs. via online tools, face-to-face vs. remote advising), (iii) whether advising responsibilities are centralized by the university or delegated to its academic departments, and (iv) whether specific advisors specialize in topics or they all are able to assist students with any program-specific and university-wide matter. The academic advising model adopted by a HEI is defined by each of these dimensions, which depend on the HEI's values, resources, and educational models.

An important aspect of any academic advising program is assisting students with course enrollment. In this process, advisors help students shape their academic path through choices oriented to the realization of their career and life goals [20, 21, 48]. The interactions between students and advisors are conducted under the assumption that, due to their experience, advisors are able to foresee the long-term consequences of the students' decisions—something that is less accessible to the latter due to their incomplete knowledge.

A long history of research suggests that course recommendation can indeed reduce the risks associated to bad enrollment decisions [18–20, 42, 54, 59]. However, when providing guidance on which courses a student should take next, academic advisors face the challenge of reconciling multiple relevant factors, such as the students' profile (e.g., academic history, interests, and skills) and the features of the courses available for enrollment (e.g., contents, difficulty, imposed workload). This makes course recommendation a time-consuming task that is also prone to errors and susceptible to personal, non-objective opinions. Logistic criteria, such as schedule restrictions and the number of places that courses can accommodate, make the task even harder. These aspects often force advisors to provide a "plan B", in case students are unable to enroll in the courses of a primary recommendation. Other challenges arise from the HEI's educational model: in universities with largely elective curricula, students have more options and thus, course recommendation tends to be easier. Less flexible curricula make enrollment decisions more critical, because poor term planning usually incurs delays that can have a long-lasting impact on the students' academic experience and success [4].

The process of course recommendation is particularly problematic when it takes place as part of an academic advising session, as these must often occur within a short time window. The problem exacerbates when students also seek assistance with personal issues that affect their academic lives. Given that advisors may serve other roles within their institution, time-related restrictions can be a barrier in achieving the goals of an academic advising system [14].

2.2 Data- and AI-based Support for Course Selection and Recommendation

The challenges explained above have motivated a plethora of tools oriented to support students and advisors in term planning and course recommendation. Most of these tools share a common denominator: they capitalize on historical data to provide users with objective information to make better, more informed decisions. At

²<https://nacada.ksu.edu>

their onset, most of these tools primarily provided data-rich environments with descriptive statistics of the courses in a program, and the progress and performance of the students (e.g., [10, 12, 13, 24, 32, 46, 49, 50, 63]). A more recent generation of tools enables users to simulate scenarios and make decisions in the light of data- and AI-based predictions (e.g., [15, 26, 52]). The benefits and drawbacks of both kinds of tools have been extensively investigated, particularly from the students' perspective. The "Cornell Experiment", for example, showed that when course grades are published online, low-performing students tend to choose leniently graded courses [7, 8]. Along the same lines, Chaturapruek et al. found that using a tool that exposes historical official transcripts and detailed end-of-course evaluation surveys lowered students' GPA—an effect that was also observed when the tool provided information about course time commitment [16].

A parallel line of research has explored how predictive tools support students' decision-making for enrollment. Backenköhler et al. proposed a tool to compose a personalized curriculum for a given student, by optimizing predicted performance and time-to-degree [4]. Jiang and colleagues describe a course recommendation tool that takes into account students' interest, prior knowledge, and development [31]. Recent efforts by Pardos et al. have focused on promoting serendipity in the exploration of enrollment options in largely elective academic programs [51]. In previous work, we investigated how a tool that presented courses' historical data and grade predictions affects students' enrollment choices [41]. We found that in the presence of grade predictions, students may be tempted to approach course selection as a grade maximization problem: they tend to choose courses with high potential GPA gain, often disregarding other important factors such as the imposed load or time commitments.

Nevertheless, when used properly, course recommendation tools provide students with a comprehensive view of their study program, and support them in making enrollment decisions in light of more relevant and personalized options [38].

While the support for students to select courses is extensive, the needs of advisors around course recommendation have received less attention. LISSA [15], eAdvisor [52], and LADA [26] are notable examples of tools that seek to support advisors in making more objective course recommendations. LISSA and eAdvisor's main goal is to support the dialogue between students and advisors through visual descriptive dashboards, following principles from the Visual Learning Analytics community [58]. LISSA also provides graduation time predictions that advisors can consider when recommending courses. The support of eAdvisor includes automatic identification of students who are unlikely to succeed in their major, enabling advisors to provide additional academic support on enrollment options. Beyond descriptive data, LADA assists academic advisors through prediction models that enable comparative analyses. When advisors build a semester plan for a given student, LADA uses clustering techniques to find similar students (according to advisee's grades and the selected courses). The tool then predicts the advisee's "chances of success" and academic risk from the data of similar students that made similar enrollment decisions at comparable points of their degree.

The use of LISSA and LADA has been studied with several samples of advisors from HEIs with varied advising models. The results of these studies show use differences between experienced and inexperienced advisors [15]. Our work builds upon—and expands—these observations by investigating how a grade prediction tool shapes advisors' strategies, and how these strategies vary according to the advisee's academic profile.

2.3 AI-supported Human Decision-making

The use of AI-based solutions for course selection has motivated a prolific body of research on how these solutions affect people's decision making and shape their attitudes, perceptions, and behavioral outcomes [35]. Aspects such as attachment, trust, fairness, transparency, and interpretability of AI-based tools are growing in importance and attention [17, 22]. These factors are also relevant for data-based tools designed to support students and advisors in course selection and recommendation. In particular, transparency in the predictions provided by automatic advising is required in order to increase confidence [26]. Some tools offer transparency by explaining the reasons behind a given prediction. Yu et al., for example, explored ways of explaining students why they had been recommended a particular set of courses using different types of explanations, each designed with a different level of personalization [62]. Other efforts have taken a more visual approach in explaining the factors that play a role in the outcome of grade prediction models for course selection and recommendation [41].

The factors that shape people's perceptions and attitudes toward AI-based tools for academic advising are more critical when such tools seek to replace the role of the advisors, mainly for repetitive tasks—such as a course recommendation. This kind of tools are described under the term "virtualized advising", defined by Thompson et al. as "*a means by which advising is provided through an impersonal means such as via an online advising program*" [56, p. 13]. Along these lines, recent work has proposed data-based tools that automatically determine the set of courses a student should take in an upcoming semester (e.g., [2, 3, 28, 43]). Some of these tools take specific designs such as chatbots (e.g., [34, 36]). Although promising, most of these tools are still at a very early stage and their benefits to real-world settings are yet to be seen. One of the main drawbacks of these tools is that their AI models can hardly incorporate factors outside a student's academic life that are, nevertheless, relevant to make appropriate enrollment decisions and recommendations. Thus, they can hardly replicate the benefits of person-to-person academic advising.

Our study surfaces some aspects of the design of AI-based tools for academic advising that may be perceived by advisors as problematic. This comprises trust in the predictions and concerns about how such tools could negatively influence the students' enrollment motivations.

3 RESEARCH CONTEXT AND RESEARCH QUESTIONS

Our research takes place within an Engineering-oriented public university with a well-established academic advising system for its 32 undergraduate programs. The study curricula of those programs

are relatively rigid, that is, students have some elective credits, but most of the courses are compulsory and cannot be avoided. Students can, though, take a course anytime they wish, provided that they have passed the course's prerequisites. In the remaining of this section, we provide details about the advising system of our research site, as well as its current technological support for the student-advisor meetings. We also describe a grade prediction tool that is part of an ongoing effort to improve the student-advisor dialogue. The eventual deployment of this tool motivates the research questions addressed in this paper.

3.1 The Academic Advising System

Our study site has a decentralized advising model in which the advising responsibilities are delegated to the university's academic departments. The departmental advisors are faculty members, who are assigned up to 40 advisees every term, depending on their workload.

All students are assigned an academic advisor—that may change throughout their student life—but advising sessions are not mandatory for everyone. Meetings with advisors are compulsory twice every term for specific groups of students: (i) first years, (ii) students with a GPA below the average GPA of their academic program, (iii) those who failed a course in the previous semester, and (iv) those who must retake a previously failed course. The first advising meeting takes place two weeks before the term starts, and is mainly devoted to address enrollment. Hence, the expected outcome of this meeting is a course set recommendation for the students' forthcoming term. More often than not, this meeting also provides information about the contents of the recommended courses and guidance on the students' time commitments and strategies to deal with the term's load. The second advising session takes place right after the midterms and is meant to monitor the students' performance and to discuss matters that may be affecting their well-being. Students who are not required to attend the advising sessions can still do so if they want to.

The academic advising system was originally designed for in-person meetings that used to take place mostly at the advisors' office. However, students could request remote and asynchronous advising sessions. The COVID-19 pandemic forced advisors and students to conduct the meetings over conference software. Before the pandemic, the vast majority of advising meetings were conducted in-person. However, nowadays, most advising sessions are conducted remotely.

3.2 Current Support

During their meetings with students, the advisors of our research site use an in-house web tool that provides access to the academic history and the records of previous advising sessions of their assigned students. The tool also presents relevant statistics (e.g., current GPA, number of failed courses, academic progress) and depicts—via line charts—the student's academic evolution, namely, passed and failed courses per academic term. This tool also allows advisors to post their time availability and students to book advising meetings.

During the term planning meetings, advisors use the tool to recommend a set of courses to students. While discussing the students'

enrollment options, advisors compose a course set by dragging textual elements that represent the advisee's available courses. Advisors can store their recommendation in the tool and retrieve it in future meetings. Students, however, do not have access to the recommendation after the meeting is over.

3.3 A Grade Prediction Tool

The university of our research site is currently taking steps to improve the dialogue between students and advisors through technology. One of these efforts is a tool named iCoRA [13], which enables the composition of course sets from those available for enrollment. By capitalizing on historical data about the different courses in the study program, iCoRA provides performance predictions for each of the recommended courses.

The development of iCoRA is motivated by a growing student population that imposes an ever-increasing load on the advisors. It also responds to the fact that when making course enrollment decisions, students are exposed to a large ecosystem of non-official—and sometimes contradicting—information about the courses they could enroll in. For instance, according to an in-house inquiry conducted at our research site, 83% of the students in the CS department ask other fellow students about the difficulty of the courses, and the reputation of the lecturers. Since the feedback obtained by the students can vary greatly depending on whom they ask, the tool counterbalances the influence of the *vox populi* by providing a more objective view that allows for more informed decisions. On the advisors' side, the tool pursues several goals: 1) providing advisors with data-based evidence to better justify their course recommendations; 2) reducing the subjective bias that advisors may introduce in their enrollment recommendations due to their learning/teaching experience (e.g., regarding the perceived difficulty of a given course); and 3) making advise more consistent across students and advisors—as the same student may receive different advice from different advisors, and the same advisor may make different enrollment recommendations to students with quite similar academic histories. Another important goal of iCoRA is to optimize the advising process, which often incurs time and financial resources from HEIs.

Figure 1 shows iCoRA's interface. The composition of course sets takes place via drag and drop operations. As an advisor drags courses from the advisee's academic program (Figure 1.a) onto the grade prediction panel (Figure 1.b), the tool updates the predicted grades for the chosen courses. Predictions are computed using machine learning models that take into account the student's performance on previous course enrollments (specially the courses' prerequisites), their success and failure history, and the chosen total workload in hours per week—based on the number of study hours of each course specified by the curriculum. The predicted grade for each course is shown as a range on a horizontal scale between 0 and 10, according to the university's grading system. This range is computed via quantile regression with gradient boosting trees. This provides us with three predictions, namely a lower bound, an average prediction, and an upper bound that we use to construct the ranges depicted by the tool. The lower and upper bounds of the individual course predictions are used to provide optimistic and pessimistic estimates for the GPA the student would obtain if the predictions became true. This GPA estimate is shown above




Figure 1: GUI of the grade prediction tool used in our study. The tool shows the courses of the student’s academic program (a). When recommending courses, advisors can drag elements from the student’s program onto the grade prediction panel (b). In response to these interactions, the tool predicts the student’s performance in each of the selected courses and computes the student’s GPA (c) that would result if the predictions became true. Clicking a course of the study program reveals its history and general information (d). The tool also shows the weekly load imposed by the composed set of courses (e).

the individual course predictions (Figure 1.c). The tool also enables access to the courses’ historical information such as the distribution of grades and the historical success and failure rates (Figure 1.d). The tool also shows the academic load imposed by a given set of courses (Figure 1.e).

To promote transparency in the prediction process, iCoRA also provides explanations for the predicted grades. These explanations convey how much the inputs to the tool’s predictive models contribute to its outputs. Advisors can access these explanations through the **Why** button associated to a grade prediction (Figure 2.a). Clicking on this button opens a modal window (Figure 2.b) showing a pie chart that depicts the relative contribution of the features used by the tools’ prediction models. These explanations are provided to make advisors aware of the factors that could impact a student’s performance in a given course. In light of these explanations, advisors could, for example, reduce the academic load imposed by a given set of courses when this feature exhibits an important negative influence on a predicted grade.

The features of iCoRA respond to the perceived needs of the university’s academic advisors. Its design is the result of an iterative design process that comprised several sketching and brainstorming meetings with academic advisors, students, as well as education, visualization, and HCI researchers. The design process also included several critique sessions in which stakeholders looked at earlier versions of the tool and provided feedback for subsequent iterations.

The effects of iCoRA on students, when they use it for term planning, have been previously investigated and reported in [41]. However, the tool’s potential benefits and drawbacks in supporting course recommendations from the advisors’ perspective are still unknown. This motivates the following research questions (RQs):

RQ1: What are the effects of grade predictions on the advisors’ approach to course recommendation?

RQ2: How does the advisee’s profile impact the advisors’ strategy to recommend courses?

The study we report below aims to answer these questions in the interest of deploying iCoRA in a real-world setting.

4 METHODOLOGY

We conducted a study to investigate the influence of iCoRA’s grade predictions on academic advisors when assisting students with different performance profiles in planning their upcoming term.

4.1 Participants

We sent email invitations to recruit academic advisors of a Computer Science (CS) undergraduate program with over 600 students. Approximately 35% of these students must attend advising sessions every term for one or more of the reasons explained in Section 3.1.




Figure 2: Clicking the *Why* button (a) associated to a grade prediction opens a modal window explaining the reasons behind the predicted grade (b). Interacting with the pie chart displays descriptions of the features and how it influences the grade.

Out of the 29 advisors of the CS program, 15 volunteered to participate in the study (8 female; 7 male; 34–59 years old; median age 41). These participants were lecturers with significant academic advising experience (median 6 years) who assist between 5 and 40 students every academic term (average 17). Most advisors ($n=11$) reported dedicating around 15 minutes per advising session. However, the advisors also remarked that the meetings’ duration vary significantly, depending on how much guidance students need.

4.2 Three Student Profiles

We asked advisors to use iCoRA to recommend courses to three fictional students who were about to enroll in their fifth semester. The students exhibited distinguishable academic performance profiles, namely low, average, and high GPA (6.50, 7.50, and 8.74, respectively). The GPA of the average-performing student was equal to that of the overall CS study program. We decided the GPA of the other student profiles based on what an *outstanding* and a *struggling* student usually mean within the program. Besides their GPA, the students had different academic histories: the high-performing student had never failed a course, while the opposite was true for the other two. The average-performing student had failed 3 courses (once each), while the low-performing student had failed a total of 17 courses (some in multiple occasions). Figure 3 shows the academic history of our three fictional students as shown in iCoRA. The red rectangles that appear behind some of the depicted course represent failed enrollment instances.

We built these three academic profiles to simulate three types of students that advisors usually encounter in their advising sessions. For each student profile, the advisors’ task was to make recommendations among the same set of 10 available courses for enrollment.

4.3 Adapting iCoRA

To address our research questions, we adapted iCoRA to provide our participants with the data they normally have at hand during the advising sessions with students. This included: the advisee’s enrollment history, their progress, and their performance evolution depicted through line charts (see Figure 4).

These modifications made iCoRA and the university’s current advising tool comparable in terms of functionality, except for three additional features: iCoRA also provides course historical information, academic performance predictions, and explanations thereof.

For all of this, iCoRA takes a visual approach. For instance, the student’s academic history as well as the courses available for enrollment are organized into a grid of visual elements with connections representing course prerequisites. This visual approach is in contrast to the current advising tool, which is mostly based on HTML form elements.

We also created a set of hardwired performance prediction models for the courses available for enrollment. The models were designed to comply with the courses’ historical difficulty, with their actual workload, and with the observation that, overall, higher workloads lead to worse academic performance among students. In particular, the relationship between workload and predicted grade is piece-wise linear with negative slopes, and is parameterized by two thresholds that define the change in the function’s slope. These thresholds correspond to the inflection points in the curve of the example shown in Figure 5. In our study, each course had three different prediction models, which accounted for the difference in expected performance for the low-, average-, and high-performing students. These models had different intercepts (the predicted grade at zero workload), different slopes, and thresholds, and were designed to never predict grades lower than 5.5 or higher than 10. The explanations for the predicted grades, available through the prediction’s *Why* button, are computed with SHAP [37], a feature-attribution explanation method based on coalitional game theory. We run the explanation module on all the inputs traditionally used by iCoRA, namely the advisee’s academic history, the selected courses’ prerequisites, and the chosen academic load.

4.4 Procedure

We conducted individual experimental sessions using video conferencing software to test and interview participants remotely. Each study session took approximately 60 minutes and consisted of the following activities:

Introduction to iCoRA. After providing consent and filling out a questionnaire about demographics and academic advising experience, each participant watched a 6-minute video that explained iCoRA’s GUI. The video described how to compose sets of courses, as well as the tool’s performance predictions and its explanations. After watching the video, participants were given the opportunity to ask questions about the tool’s functionality.




Figure 3: Academic history of the student profiles used in our study. Red rectangles behind the courses of the study program represent past failed enrollment instances.

Course recommendation tasks. Advisors were then asked to use our adapted version of iCoRA to recommend a course set for the upcoming academic term to each of our three fictional students. Advisors were told that these advisees would have online access to the course recommendations made during the study session. No restrictions were specified about the number of courses advisors were allowed to recommend or about the time they could spend in the task. We shuffled the students' order across participants to reduce order effects.

Closing Interview. The experiment concluded with a semi-structured interview in which we asked participants about their perspectives on iCoRA. This included their general opinion on the tool, perceived benefits and drawbacks, and their thoughts on an eventual deployment of the tool at their institution.

4.5 Data Collection, Statistical Tests, and Qualitative Analysis

We recorded the courses that each student was advised to take, including any partial sets that the advisors built with iCoRA while making their recommendations. We characterize these course sets by their incurred workload and GPA gains. We also recorded how frequently and how long advisors engaged with the **Why** button's explanations of the predicted grades.

To verify if the student's profile had an impact on the nature of the recommendations, we conducted a Kruskal-Wallis H test among the three groups of observations drawn from each student profile (i.e., course set's workloads, GPA gains, and engagement with the explanations). The null hypothesis of the Kruskal-Wallis H test is that *the median of the population is the same, regardless of the student*




Figure 4: General information GUI component of the modified version of iCoRA used in our study. This interface component provides advisors with information on the student’s enrollment history and performance evolution.

profile—otherwise it is said that some profiles stochastically dominate others. When the test revealed stochastic dominance (with a significance level $p < 0.05$) from at least one of the observation groups, we conducted a post-hoc pairwise Dunn’s test with Bonferroni correction to detect the pairwise dominance relationships between the profiles.

We also captured 312 minutes of the participants’ advising process in the form of video screen captures. We audio recorded the closing interviews (181 minutes in total) and fully transcribed them. Using an inductive, thematic analysis technique [11], we structured

the participants’ remarks from the interviews. Initially, two authors coded at least seven interviews independently. This initial coding sought to identify emerging patterns in the advisors’ strategies and rationale for course recommendation, as well as their vision and the perceived benefits of the tool. In a second stage, we redistributed coding assignments so that each interview was analyzed by at least two different researchers. Through regular meetings, the research team iteratively revised the initial topics and refined higher-level themes that emerged from the data.




Figure 5: Predicted GPA as a function of the workload for the *Software Engineering II* course. In our study, for all the courses available for enrollment, the predicted loss in GPA points due to an additional hour of workload depended on whether the advisee had a low, average, or high overall workload.

Student performance	No. courses			Total hours			GPA gain Γ^s		
	Median	Mean	Std. dev	Median	Mean	Std. dev	Median	Mean	Std. dev
<i>high</i>	4	4.27	0.96	33	34.0	7.14	-0.03	-0.02	0.03
<i>average</i>	4	4.13	0.63	33	33.4	5.30	-0.003	-0.001	0.01
<i>low</i>	4	3.80	0.67	33	30.4	6.60	0.06	0.06	0.04

Table 1: Number of courses and total hours recommended per student profile.

5 FINDINGS

This section presents the findings of our study with the academic advisors. We begin describing the quantitative results derived from the advisors' decisions and interactions with iCoRA. This is followed by the results of our thematic analysis of the interviews.

5.1 Quantitative Analyses

We first describe and compare the advisors' recommendations for each of the student profiles of our experiment. This comprises (a) an analysis of the academic load of the recommendations for each profile operationalized by the total number of study hours per week, the number of courses, and the predicted GPA gain; and (b) a visualization of frequent combinations of courses recommended by the advisors. In a second stage, we analyze (c) the time spent on the recommendation task, (d) the total number of drag-and-drop steps carried out before making a recommendation, and (e) the time invested in reading the tool's explanations for the predicted grades.

5.1.1 Advisors' Recommended Load. A natural way to characterize the recommendations made by the advisors is to compute the study workload incurred by those recommendations. This can be operationalized by the number of recommended courses and the total number of study hours (per week) of those courses. We did not consider the grades and GPA predicted by iCoRA in this analysis because, due to our experimental protocol, the prediction models across courses and student profiles had different difficulty levels: they had a different intercept (i.e., predicted grade at zero workload) and a different slope (i.e., the models penalized the grade differently for each additional hour of work). This makes absolute GPA comparisons meaningless. That said, we can indeed compare predicted GPAs among student profiles if we control for the difficulty of the prediction models used by a given advisor. To do so, we report the *corrected gain in GPA* of a recommended set of courses, denoted by Γ . Given a set of courses C_s recommended by a given advisor to a student s , the corrected gain in GPA is computed as

$$\Gamma^s = \frac{\gamma^{C_s}}{\gamma^{C_{avg}}} (\overline{gpa}_{opt}^{C_s} - gpa_s),$$

where:

gpa_s is the advisee's initial GPA (6.50, 7.50 or 8.74 for the low-, average-, and high-performing student respectively),

$\overline{gpa}_{opt}^{C_s}$ is the GPA's optimistic prediction³ (see Section 3.3) if the student took the courses in C_s , and

$\gamma^{C_s}/\gamma^{C_{avg}}$ is the correction factor.

The term $\gamma^{C_s} < 0$ is a measure of the difficulty of C_s defined as the average loss in GPA points for every hour of additional workload for the courses in C_s . This value is computed as the average slope of the linear models used to predict each of the grades in C_s given the total number of study hours. To account for the fact that γ^{C_s} is different across the different student profiles, we divide it by $\gamma^{C_{avg}}$, the average difficulty of the models used by the same advisor for the average-performing student. That way, we make the GPA gains comparable across student profiles for the same advisor.

The results of this study are reported in Table 1, where we can observe that on average, advisors assigned lower workloads to the low-performing student. That said, we could not reject the null hypothesis of the Kruskal-Wallis H test for the total hours and for the number of courses. We can also see that the GPA gain is positive only for the low-performing student, i.e., the advisors managed to increase the GPA of the low-performing student, whereas for the other profiles the recommended set led to a (very) slight performance drop. This is confirmed by the Kruskal-Wallis H test that reveals statistical dominance ($H(32) = 32.70, p < 10^{-7}$). The Dunn's posthoc test, reported in Table 2, shows a significant difference between the low-performing student and the other advisees. This suggests that the advisors not only focused on protecting the low-performing student from high workloads but also took more care of improving their GPA. Similar trends can be observed when looking at the non-corrected GPA gain $\overline{gpa}_{opt}^{C_s} - gpa_s$ (although the values are not the same). These results are in line with the comments of some of the advisors who were rather optimistic about the grades of the high-performing student despite the predicted decrease in the advisee's GPA (see Section 5.2).

5.1.2 Frequently Recommended Courses and Courses Combinations. To identify "popular formulas" within the advisors' recommendations, we conducted an analysis based on course frequency, itemset mining, and agreement. Figure 6 shows a frequency ranking of the different recommended courses per student profile. The figure reveals that the course *Entrepreneurship and Innovation* was overall the most frequently recommended—10 out of 15 times for the average performing student, 11 times for the other students. This may be explained by the fact that this course is, albeit time-consuming,

	<i>low</i>	<i>average</i>	<i>high</i>
<i>low</i>	1	0.000901	4.96×10^{-8}
<i>average</i>	-	1	0.13
<i>high</i>	-	-	1

Table 2: Post-hoc comparisons for the GPA gain. Corrected p-values shown.

³The gains for the GPA's pessimistic prediction are always negative.

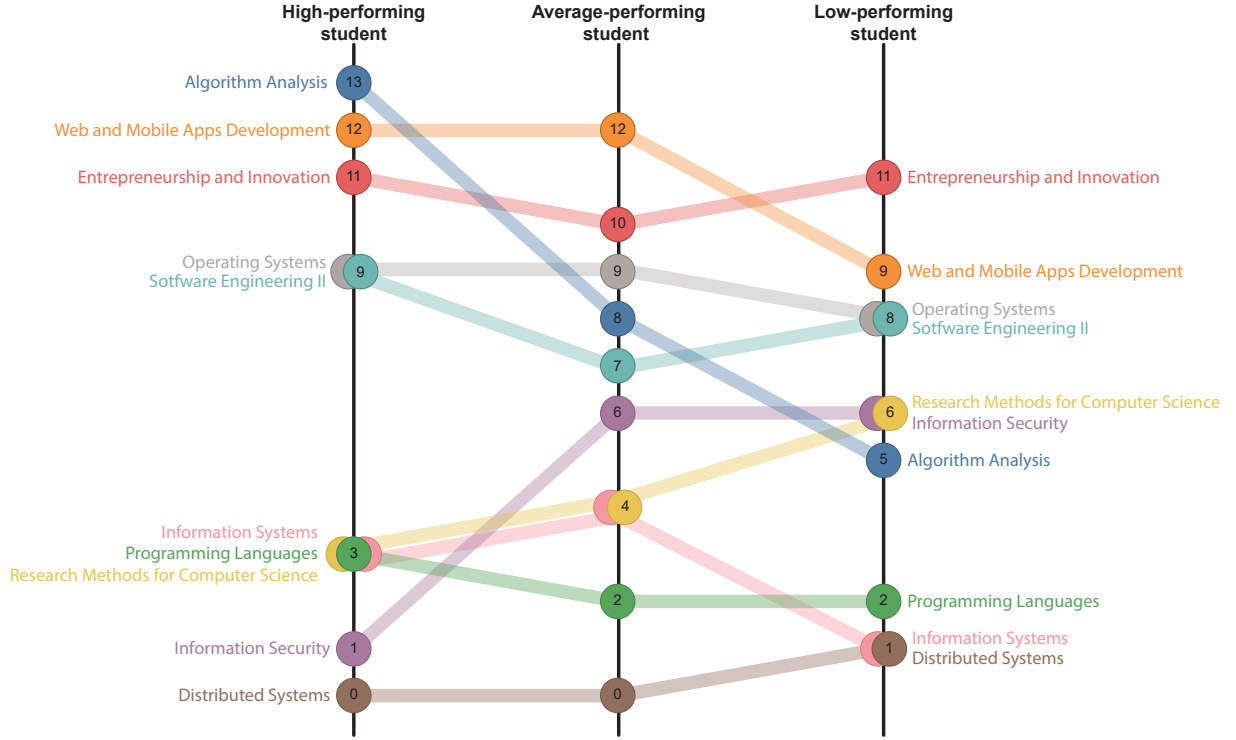


Figure 6: Courses' recommendation occurrence. The number within the circle representing each course indicates the number of total times it was suggested by the advisors to the corresponding students.

usually perceived as easy for CS students. This makes it a good bet for advisors. The *Web and Mobile Apps Development* course was also a popular choice: it was recommended 12 times to the average- and high-performing students, and 9 times to the low-performing student. This makes it the most recommended course for the average-performing student, and the second most popular suggestion for the other students. This probably has its roots on the course's perceived practical value. One of the advisors defined it as “a door to the job market” [P15] in reference to the observation that students can capitalize on the skills acquired immediately after passing this course.

Despite some general trends in the selection of the courses, we also observed trends that depend on the student's profile. For instance, the *Algorithm Analysis* course ranked the most popular in the recommendations for the high-performing student (13 out of 15 times), whereas it was recommended only 5 times to the low-performing student—always in combination with *Web and Mobile Apps Development*. As stated by one of the advisors, this course is theoretical and requires students to work hard. Conversely, the *Information Security* course was recommended 6 times to all students except for the high-performing one, who got it recommended only once. Advisors may have not prioritized this course as it is not the prerequisite of any other.

Finally, we conducted an agreement analysis based on the Jaccard coefficient. If C_1 and C_2 are two sets of courses recommended to the same student by two different advisors, then the Jaccard

coefficient $\mathcal{J}(C_1, C_2)$ quantifies the consensus between those recommendations. We compute $\mathcal{J}(C_1, C_2)$ as the number of courses recommended by both advisors divided by the joint number of recommended courses, put differently:

$$\mathcal{J}(C_1, C_2) = \frac{|C_1 \cap C_2|}{|C_1 \cup C_2|}$$

Identical recommendations yield a Jaccard score of 1.0, whereas values closer to 0 denote high disagreement. We computed $\mathcal{J}(C_1, C_2)$ for every pair of recommendation sets leading to $15 \times 14 / 2 = 105$ pairs of recommendations (recall that the score is commutative) per student profile. We report the average Jaccard score in Table 3. The results suggest a trend towards disagreement for the low-performing student, i.e., $\mathcal{J}(C_1, C_2) = 0.34$, whereas, for the high-performing student, each pair of advisors agreed, on average, almost on half of the courses that would be suggested to the student. In all cases, the standard deviation is significant and almost constant.

High		Average		Low	
Mean	Std. dev	Mean	Std. dev	Mean	Std. dev
0.48	0.20	0.37	0.21	0.34	0.22

Table 3: Recommendation agreement (via the Jaccard coefficient) for the different student performance profiles.

Student performance	Time (minutes)			No. of actions		
	Median	Mean	Std. dev	Median	Mean	Std. dev
<i>high</i>	6.36	6.55	3.91	35	30.27	18.07
<i>average</i>	6.04	6.53	4.56	35	36.06	23.58
<i>low</i>	8.20	9.64	5.17	35	41.27	28.05

Table 4: Time spent and number of actions (addition of removal of courses) before submitting a recommendation.

5.1.3 Interaction Effort of the Recommendation Task. We also analyzed the effort invested by the advisors in the recommendation for each student profile. We characterize the advisors’ effort in terms of *a*) the time they interacted with iCoRA before submitting a course set recommendation, and *b*) the number of drag and drop operations this interaction incurred. Table 4 reports the average values of these metrics across all advisors. The values suggest that, on average, the advisors spent more time making a recommendation for the low-performing student—which may explain why they tended to disagree more on the courses this student should take. A similar trend can be observed for the number of actions—addition and removal of courses via drag-and-drop operations—carried out during the advising session, so to say, that the less performing the student is, the more effortful the recommendation tends to be. Despite these results, we could not reject the null hypothesis of the Kruskal-Wallis H test, very likely due to the high standard deviation of these observations.

5.1.4 Interest in Explanations. We configured iCoRA to log the advisors’ interactions with the explanations of the predicted grades for the individual courses. While almost all advisors (14/15) opened the grade explanations at some point during the recommendation tasks, overall, they showed little interest in this functionality. We registered an average interaction time of 1.28 minutes per advisor (with a standard deviation of 1.02 minutes) during the entire experiment—if we exclude the advisor who did not invoke the explanations at all. On average, advisors invoked the explanations 3.16 times during the entire experiment. However, most interactions happened only for one of the three advising tasks. If we consider each of the 45 recommendation tasks, only 19 tasks incurred an interaction with the explanations, with those making an average time of 55.33 seconds (with a standard deviation of 52.78 seconds). Out of those 19 recommendation tasks, 11 saw more than one invocation of the explanations. We remark, however, the high dispersion of the time observations as they span from 10 to 179 seconds. Five interactions lasted more than one minute, and two of them are explained by the fact that the advisors kept the explanations open while interacting with the experimenter. In the other three cases, the advisors took the time to understand the explanations and thought out loud. Finally, we also notice that the low-performing student sparked

more interest in the explanations, as shown in Table 5. This goes in line with the longer recommendation time invested by the advisors when dealing with this student. That said, these results must be taken with a grain of salt given the high variance exhibited by our observations.

5.2 Qualitative Analysis

We present in this section the findings of our qualitative analysis of the interview data and the video recordings of the advisors’ interactions with iCoRA. These observations provide a broader interpretation of the quantitative results presented earlier, as they uncover subjective factors that advisors brought into their decision process and their perspectives on iCoRA.

5.2.1 General Perspectives and Perceived Benefits. With no exception, all advisors were enthusiastic about the perspective of having at hand a tool like iCoRA. Some explicitly stated appreciating its functionality in comparison to the tool they currently use to conduct the advising sessions: “[iCoRA] could be extremely useful because it goes one step further than what currently exists [for the advising sessions]” [P11]; “There was information that I did not know before and now, because of the tool, I do. So, the tool helps me contrast the ideas I had in mind about certain courses of our curriculum.” [P15]. The most recognized advantages were centered around the data aspects of the tool: “It gives us, advisors, statistics and historical data to make better decisions” [P02]; “Because of the data it provides, this tool could make students understand more clearly about whether or not they could fail a course.” [P03]. In particular, the tool’s grade predictions were seen as instruments to provide students with more objective recommendations: “Making blind decisions never has advantages. I think with a tool like this, students will get a more ‘educated’ recommendation” [P09]; “Based on the data, I can explain to the student, ‘this is the probability that you will fail’. I can provide advice based on real data [...] to convince the students [which courses are appropriate to take]” [P10]; “This tool would allow me to provide more well-founded recommendations. In turn, students will receive more accurate information based on data and not only on the experience of the advisor” [P09]; “The data used by the tool makes the advice we can provide slightly stronger” [P03].

No. of tasks				Average Frequency				Average Time (seconds)			
<i>total</i>	<i>high</i>	<i>avg.</i>	<i>low</i>	<i>total</i>	<i>high</i>	<i>avg.</i>	<i>low</i>	<i>total</i>	<i>high</i>	<i>avg.</i>	<i>low</i>
19/45	5	6	8	2.28	2.00	2.4	2.38	55.33 (52.78)	42.25 (35.89)	42.26 (31.13)	71.38 (69.33)

Table 5: Advisors’ engagement with the explanations per recommendation task and student profile. Average values are provided with the standard deviation in parentheses for the times.

5.2.2 Advisors' approach to recommending courses. Our video and interview analyses revealed that advisors approached course recommendation mainly by considering their experience and personal perspectives on the courses. All advisors looked at the grades predicted by iCoRA at some point during the study session, but the predictions were not taken into account in all the course recommendations tasks.

When looking at the predictions, advisors considered both the grade of each selected course and the potential effect of the courses on the student's GPA. The latter was the main criterion in at least one course recommendation for 13 advisors: *"I was trying that their GPA did not decrease; I did not really pay attention to the courses' individual grades."* [P10]; *"I left the 'Algorithm Analysis' course out because it was going to decrease the students' GPA more significantly. So, I decided to replace it with the 'Entrepreneurship and Innovation' course."* [P05]; *"I went for the course sets that gave the student the highest GPA gain. The courses' individual grades the tool predicted didn't help me much [to make a decision]. My goal was that the student improved their GPA."* [P08]. In all these instances advisors described this GPA-driven strategy as their usual approach to course recommendation.

When prompted on which other factors they considered, most advisors (12/15) mentioned the order the courses appear in the study program. As stated earlier, the curriculum we worked with is relatively inflexible (in contrast to many Global North universities, that include a large number of elective credits). Hence, students do not really have much room for optional courses. Also, as courses have set prerequisites, there are few paths students can take at a given point of their academic history. In such a context, advisors considered course order an important criterion to decide which courses recommend: *"I consider the logical relationship between courses and their order in the curriculum."* [P11]; *"I tried to make them take the courses in an ordered fashion. In my opinion, the courses that are not prerequisites of others could wait."* [P10].

Regardless of whether they closely observed the predicted grades, the impact on the students' GPA, or the order of courses in the curriculum, all advisors tried to balance the load of the sets they suggested: *"I would suggest them to take two hard courses and two easy ones."* [P11]. Balancing the term's academic load was particularly important when making course recommendations to the low-performing student: *"Because of the GPA this student has, I would go for not recommending them a course set with a high load."* [P13]; *"When they have a GPA like this, one normally does not recommend them to take more than four courses [per term]."* [P08]. As our quantitative results show, advisors spent on average more time trying to assemble a course set for this student. In most cases, this was due to repeated attempts to find a set with minimal load and low negative impact on the student's GPA. One advisor even tried to find "easy" courses to counterbalance the load imposed by science courses: *"None of these courses is easier than 'Entrepreneurship and Innovation'! [...] And there are no more humanities courses⁴ available for this student to take!"* [P10]. Advisors also based their recommendation on the skills required to pass the courses: *"Besides the general academic history, I pay attention to how they [students] have*

done in Computer Science courses. In particular, in courses that require programming skills [...] If they have performed poorly in those courses, I would exclude similar ones from my recommendation." [P15].

The strategies all advisors took to recommend courses to the low-performing student are in stark contrast to those applied for the other student profiles of our experiment. In the case of the high-performing student, for example, most advisors were willing to recommend high-load sets even in the presence of not-too-optimistic predictions: *"It's going to be a complicated term but since this is a good student, I'd also recommend them the 'Information Security' course—despite the fact that it may decrease their GPA because it will be a busy term."* [P10]; *"I paid attention to the GPA of all the students. However, there was one student who had a very good GPA, so I recommended them to take the entire row of courses [as they appear in the visual depiction of the curriculum]."* [P13]; *"I always tried to keep a healthy balance between performance and academic load. However, because this was a good student, there is a guarantee that they will be able to get through their [high-load] term."* [P09].

Our inquiries also show that advisors relied heavily on their experience and knowledge of the curriculum's courses. 9/15 advisors overlooked or decided to ignore the courses' historical and official workload information. These advisors suggested that their experience was enough to have a clear idea of the difficulty level of each course: *"From experience, one already knows which are the courses in which students usually have problems."* [P13]; *"This [the course's number of hours] is what the curriculum specifies, but not what the students tell me [...] Personally, I would ignore this information."* [P15]. Four advisors relied so much in their experience that they completely ignored the predictions of the tool and used it only to compose sets and send their recommendations to the students: *"I would make my recommendations based on my experience, because [as advisors] we know the stuff."* [P12]. In these cases, advisors reviewed the student's academic history thoroughly to get a sense of the student's academic weaknesses and strengths. This information was also important to establish the students' affinity to specific types of courses and to identify course combinations that proved problematic in the past: *"When there are failures, I would like to know which courses the student passed and which ones they failed. For me, it is not enough to know how many they passed. I need to know if the problematic courses came from one branch [of the curriculum] or another, or if the student needs to take a course for a second or third time."* [P12]. Two advisors used the student's academic history to know how many courses the student took in previous semesters and adapt the load accordingly: *"The fact that the student took only 4 courses [in the previous term] suggests that they reduced the number of courses either due to the pandemic, to family issues, or to work duties [...] Based on that and assuming this situation will not change, I would recommend them between four or five courses."* [P07].

Our analysis also reveals that advisors had an already well-established idea of the relative importance of the courses. Importance was often defined in terms of the contents covered in a course and their applicability in practical matters of a student's life. For example, some courses were considered important for the students who, after that academic term, will be applying for internships or jobs. Such criteria prevailed, even if we consider that students are not obliged to take the courses in a specific order, and even when

⁴As the study was conducted in an Engineering-oriented university, both students and advisors often perceive courses from the humanities and social sciences as easier.

the tool presented pessimistic predictions: “*This student must take the ‘Web and Mobile Apps Development’ course. It doesn’t matter that it has a low grade predicted in both course sets.*” [P10]; “*Web Development is one of the most useful courses for students*” [P01]. Besides the short-term professional perspectives the *Web and Mobile Apps Development* course offers, this course is also a co-requisite⁵ of *Software Engineering II*. This type of relationship between courses also influenced their perceived importance: “*Also, I find it important that the student takes the ‘Software Engineering II’ course [this term] because otherwise, they will be in trouble.*” [P01]. The latter quote makes reference to the fact that the courses *Software Engineering I* and *II* have a common group final project, hence it is important that members of the same group take the courses at the same time.

We saw some level of iteration during the composition of the course sets. That is, in eight instances, advisors composed more than one set of courses and assessed their advantages and disadvantages for the students in light of the tool’s predictions. Having said that, we observed little engagement with explanations and with the courses’ historical information. Advisors trusted their guts when it came to characterizing a course, regardless of how the tool presented it. They considered some courses easier than others but not because of the information provided by the tool. Rather, it seems like they had a preconceived opinion on which courses were easier or harder than others. Also, the workload was mostly observed in terms of the number of courses, not the number of hours a student would have to invest.

5.2.3 Advisors’ Concerns. Albeit largely enthusiastic, our participants’ responses to the tool also raised concerns on the implications of adopting iCoRA as part of the university’s advising system. The main concerns revolved around how reliable the tool’s predictions were and how much they could (and should) be trusted: “*It is impossible to know if the predictions will come true or not next term.*” [P09]; “*I do not know the certainty that the prediction will be fulfilled.*” [P10]. In a few cases, advisors stated not to trust the predictions at all: “*I can’t really trust this [the tool’s predictions]. I just don’t think the grades of a group of students can be used to predict the performance of another.*” [P11]. That said, our video analysis revealed that most advisors assumed that the predictions could be trusted to some extent and that at least for the sake of the experiment, the explanations were consistent with the participants’ experience and expectations. These comments, however, suggest that an eventual deployment of iCoRA (or a similar tool) in a real-life setting, will require additional actions and mechanisms to generate trust. One of those mechanisms are iCoRA’s explanations, to which advisors paid little attention as explained earlier.

One advisor proposed checking whether the students followed the recommendations as a way to evaluate the impact of the tool. The advisor’s rationale was that monitoring how closely students followed the counselor’s advice could be seen as a proxy of how good the tool is in fulfilling its goal: “*We should get some sort of feedback on whether or not the student followed the advisor’s recommendation. Students should perhaps state [through the tool] why they decided (or not) to choose the courses that were recommended and the tool could report that to us.*” [P12].

⁵If A is co-requisite of B, then A cannot be taken after B. They can still be taken at the same time.

Some advisors, who advocated for making iCoRA available to the students, also expressed concerns on the implications of doing so. In particular, they reflected on the undesired effects of grade predictions on the students’ behaviors and motivations at enrollment time: “*The tool could push the students to always focus on their grades and not on what they are going to learn [in a course], as it is purely based on the grades but provides no information about the contents of a course.*” [P10]. To avoid this, one advisor suggested showing the probability of passing or failing the courses, rather than predicting grades: “*I would rather see the probability of the student passing the courses. If there is a 15% chance that the student will pass a course, the tool would show that in red. On the other hand, if the probability is high, over 70% for example, everything would appear in solid green.*” [P01].

The presence of the advisor was seen as paramount to mitigate the potential drawbacks of the tool: “*Some students could misuse the tool, especially those who would not know how to interpret what they see. To avoid this, the presence of the advisor is crucial, especially for very young students and for those with no affinity to statistics.*” [P04]. Along these lines, others also highlighted that the presence of the advisor is needed due to the aspects that any tool or predictive algorithm cannot deal with: “*[A tool like this] should not be a replacement of the advisor, because we do more than just recommending courses. There are human aspects that only the advisor can address during the advising sessions.*” [P04]; “*Obviously, we want to understand trends and the tool provides a lot of information for that [...] However, there is no qualitative information [in iCoRA], only quantitative data. The student’s personal and emotional problems are not captured by the tool.*” [P15].

Finally, advisors reflected on how the tool would change the duration of their advising sessions. Rather than seeing iCoRA as a tool to make the advising sessions shorter, the general opinion was the tool could indeed make the advising sessions longer, because it could motivate advisors to play around with the grade predictions until finding a course set that satisfies them: “*By having the ability to compose several sets, one could spend a lot of time trying to find suitable combinations for the student.*” [P13]; “*Advisors could use the tool to design a plan before the student comes to the advising session. But either way, this implies spending more time.*” [P10]; “*It may take longer to carry out the advising session because there is not a unique path of interaction in the tool.*” [P04].

5.2.4 Supporting the Dialogue Between Advisors and Students. Advisors’ statements during the interviews suggest that they deem their role in supporting students’ enrollment important. Regardless of the students’ academic history, advisors conducted the advising process carefully, were always willing to provide relevant and effective course recommendations that benefit students, and often reflected aloud about the potential impact of their recommendations in the light of the student’s circumstances.

Despite all the effort put into producing recommendations, all advisors, with no exceptions, acknowledged the importance of letting students have the final saying when it comes to selecting courses. For this to happen more effectively, advisors suggested that students should be able to explore the courses’ historical data and have access to the tool’s predictions before the advising meetings take place: “*It would be great if students could come to the advising*

sessions with an input after having used the tool.” [P12]. Rather than delegating the full responsibility of choosing good courses to the students, advisors proposed deploying the tool as a way to make the student-advisor more effective and efficient: “If the tool was available to both students and teachers [advisors], the students could compose and bring [to the advising sessions] their own course sets. This would speed up our dialogue.” [P08]; “It would be great if the student could fill in [in the tool] everything before the advising meeting, so that [through the tool] they inform advisors which course sets they are interested in.” [P01]; “Students could benefit from a tool like this to explore [their enrollment options] on their own. If they come [to the advising meetings] with a scenario they have built in advance, that could decrease the time needed for the advising sessions.” [P04]. Along these lines, some advisors also acknowledged the need for mechanisms to motivate students to play with the tool: “We could encourage the students to use the tool by making its use a prerequisite for the advising sessions.” [P14].

Some advisors suggested a mixed approach in which the tool provides an initial recommendation that can be later discussed by both parties during the advising meetings: “The tool could say ‘this is the most favorable scenario’ and it should be possible for us [students and advisors] to customize that initial recommendation [during the advising sessions].” [P04]. Along the same lines, two advisors suggested that the tool should enable the comparison of different course sets to assess the potential effects of these recommendations in light of their differences: “It seems to me that we should be able to quickly compare and see the differences between sets. You should be able to visualize the differences between the sets one and two.” [P01]; “It would be nice if we could see them [the course sets] all at the same time, like the comparisons of items one can do when shopping online, in which it is possible to identify differences among them.” [P04].

All these comments highlight iCoRA’s potential to make the advisor-advisee dialogue more effective, without replacing the advisors and without overriding the students’ interests in, and affinities to, specific courses.

6 DISCUSSION

Our discussion is guided by the research questions formulated in Section 3, about the effects of grade predictions (**RQ1**) and the advisee’s profile (**RQ2**) on the advisors’ recommendation approach. We build upon the quantitative and qualitative analyses presented in the previous section and structure our analysis in terms of the advisors’ strategy to make recommendations and their attitude towards iCoRA. Last but not least, we discuss the ethical considerations and implications of using grade-based predictions for assisting term planning.

6.1 Advisors’ Approach

In regards to **RQ1**, the qualitative analysis presented in Section 5.2.2 suggests that, in general, the advisors’ opinion and experience prevailed over the predictions made by iCoRA. A recurrent behavior during the experimental advising sessions consisted of focusing attention on a course and pondering its effect on the student’s performance based on the advisors’ knowledge of the study program. Such knowledge comprises three aspects: (a) the advisors’ previous exchanges with the students, (b) the advisors’ perception

of the actual workload of the courses, and (c) the general institutional enrollment guidelines. Thanks to their continuous contact with the students, advisors count on an informed view of which courses are generally challenging. Advisors reconcile that information with what they know about the courses (e.g., contents, theory vs. practical components, presence of group projects, professional perspectives), as well as with their intuition about appropriate enrollment choices. In general, advisors observed on two enrollment guidelines: follow the order of the courses in the study program whenever possible⁶, and balance performance and workload. In many cases, advisors’ prior views not only overrode iCoRA’s predictions, but also the course’s workload and time commitment information the tool provided. That is, even if a course incurred many hours of work according to the study program, advisors overlooked that information in favor of what they knew about the course’s actual demands. This preponderance of the advisors’ own view over the tool’s predictions stands in sharp contrast to what has been observed for students, who based their decisions mostly on the information provided by the tool, including both the workload and predicted grades [41].

While advisors mostly trusted their intuition, we should not conclude that they completely overlooked the information provided by the tool. In eight cases, the advisors thoroughly examined the students’ academic history to build a more precise profile of the advisee. By looking at how many (and which) courses the students had taken the previous semester or by inspecting their performance in some key courses (e.g., courses that require programming skills), advisors built a more specific picture of the student’s performance, preferences, and potential circumstances. This aimed to provide personalized advice. Two advisors (P01, P14) hypothesized that one of the students might be working, or going through family issues, from the fact that they had taken few courses in the previous semester. The advisors thus used that information to reduce the workload incurred by their recommendations. While this deductive exercise is presumably a consequence of using fictional students in our experiment, it does confirm that the advisors do consider the students’ unique circumstances when recommending courses. By circumstances we mean the students’ personal situation, their potential preferences for certain courses, or their preference for particular enrollment strategies such as taking the courses in order, taking a specific number of courses per term, or minimizing the number of semesters.

But the advisors did not only use iCoRA to overcome the lack of explicit personal circumstances. Our evidence suggests that they did look at the predictions made by the tool as a sort of validation or fine-tuning of their intuition. This validation goes in line with their internalized tendency to preserve the student’s GPA—in continuous trade-off with the workload. Some advisors would look at the predicted GPA every time they added a course to the recommendation, whereas others would first apply their personal strategy to compose a recommendation and make small adjustments at the end based on the predicted GPA. Regardless of the recommendation strategy, the advisors based the fine-tuning mostly on the predicted GPA, and looked rarely at the individual grade predictions for the courses.

⁶This excludes prerequisites since they are enforced by the structure of the curriculum.

Three advisors (P01, P07, P15) who looked at the historical information of the courses (difficulty, approval rate, and workload) mostly ignored this information in favor of their own opinion about the courses. The remaining advisors completely overlooked the course information component of the tool.

Concerning **RQ2**, there is a clear distinction between the strategy applied to advise the low-performing advisee and the other students. Advisors were more willing to take risks for the advisees with better performance profiles because they believed these students could overcome pessimistic grade predictions. This is supported not only by the advisors' comments but also by our quantitative analysis: The GPA gains for the low-performing student were positive in contrast to the other students. While we could not show statistical significance for the differences in advising time among the different profiles, we did observe higher averages and dispersion for the time invested in the low-performing student. This profile also incurred the highest disagreement among advisors, as shown in Table 3.

6.2 Advisors' Attitude and Concerns

Our observations indicate that advisors had an overall positive attitude towards iCoRA. Four of them (P01, P02, P13, and P15) praised the fact that the course recommendation revolves around a visual representation of the curriculum, contrary to the current advising tool. This is particularly important (for both advisors and students) because the study program of our research site is highly structured and restricts the courses the students can take at a given time. Advisors also appreciated the data-supported advising paradigm implemented by iCoRA, which they deem to be a good asset for improving the quality of their course recommendations.

Most advisors considered that having access to the historical information of the courses and the students' academic performance allows for more informed decisions that would enrich the value of the advising meeting. In that regard, they all agreed that a tool like iCoRA is not meant to replace the advisors, but rather empower stakeholders with pertinent, more objective information beyond the *vox populi*. All participants insisted on the importance of the advising sessions for students to discuss extra-curricular aspects that nevertheless have an impact on their performance and welfare.

Despite the overall positive attitude towards iCoRA, the advisors expressed some concerns about the deployment of such a tool to support the advising sessions. The first source of concern was the trustworthiness of the predictions. In most cases, the advisors did not question the validity of iCoRA's verdicts and assumed it was worth looking at the grade predictions for validating their recommendations. However, two advisors (P09 and P11) did not trust the predictions and argued that such functionality would require extensive validation before deployment. Besides, it is not clear whether the bulk of the advisors did not challenge the predictions because these were given in the context of an experiment—which may also explain the little engagement with the explanations. All this implies that any eventual deployment of the tool will require additional work to guarantee users' trust.

A few advisors pointed out that the panel containing the historical information of the courses was overloaded and could be simplified by removing some of its elements. This can explain why most participants took no notice of this information and relied

on their own knowledge of the study program. Finally the general consensus among advisors is that using iCoRA without prior preparation, would not reduce the duration of the advising sessions because, as we saw for the low-performing student, it encourages the advisors to invest more time and effort to find a well balanced set of courses. That said, such a phenomenon can be easily mitigated if at least one of the actors, either the advisor or the student, prepared one or two possible scenarios to seed the discussion.

6.3 Advisors' Vision on a Grade Prediction Tool

As explained in Section 3.3, the design of iCoRA is the result of an iterative process that involved different stakeholders, including academic advisors [13]. None of our participants, however, took part in the tool's design process. Thus, our experiment was also a first-time introduction of the tool to some of its future users. This allowed us to capture the advisors' vision of how a tool that provides grade predictions for course recommendation should look like. We discuss here the implications of the functionalities our participants mentioned during the interviews.

If the tool were to be used by students, including information on the courses' content may bring an important benefit: It would provide students with a fuller picture of their enrollment options and may channel part of the students' focus from the predicted grades toward the courses' learning outcomes. A similar effect could be achieved on the advisors' side by replacing the grade predictions with probabilities of success (or failure). This design could perhaps shift the advisors' attention away from the student's GPA, in favor of a less performance-driven course recommendation strategy.

The vision of other advisors included logistical aspects of the enrollment process. More specifically, they suggested augmenting iCoRA with the time schedules of the courses. That would make sure that the outcome of the advising session is a feasible recommendation. Otherwise, students face the risk of having to alter their original plan during the enrollment period or have a backup plan in case the recommended course schedule is not feasible. Logistical aspects were also mentioned regarding the organization of the academic advising sessions. Some advisors suggested that to have a starting point for the discussion, the students should have access to the tool before meeting their advisors. Others indicated that the tool could automatically propose sets of courses to the advisor—sets of reasonable workloads that, at the same time, optimize for GPA. This functionality, however, does not take into account the student's preferences and skills and would require integrating the courses' timetables.

Finally, a few advisors advocated for comparisons between course sets in iCoRA to contrast the pros and cons of two (or more) recommendation strategies (e.g., one conservative versus one more challenging), and to provide alternative recommendations to the students—in case they change their minds or run into schedule constraints. In this respect, the comparison could also include course sets prepared by the students before the advising meeting. This would require making iCoRA accessible to the students, and could potentially speed up the advisor-student dialogue.

6.4 Ethical Considerations of Assisting Term Planning with Grade Predictions

During advising meetings for term planning, academic advisors can provide the scaffolding that students need to make better, more informed, and more objective enrollment decisions. Our current research direction is to support this process via grade-based predictions. Grades are meant to convey important information to a wide range of stakeholders [44]. In the market context of our research site, for example, the students' GPA is carefully considered by potential employers during applicant screening. Using grade-based predictions also has a practical value: given that GPA data is readily available in most HEIs, it can be *easily* used in the deployment of AI-based tools oriented to assist advisors in recommending courses. However, assisting course recommendation solely through GPA-based predictions could be detrimental to the mission of any academic support system. As with any AI system, grade prediction models may contain biases (e.g., demographic, socioeconomic, ethnic) and may reflect discriminatory behavior toward certain groups [30, 40]. Another problematic aspect of grade-based predictions is that they are insufficient to holistically describe a student's performance [55]. The student's grades cannot reflect their development of learning outcomes, levels of engagement, or learning style [1, 47, 53]. Finally, grade-based prediction models are insufficient to capture the students' circumstances beyond their academic performance. When considering the students' enrollment options, there are factors beyond their grades that play a role. Aspects such as the students' life goals, extracurricular activities, and family or health circumstances must also be considered when advising students on what courses to take next. Making course recommendations while ignoring aspects beyond predicted grades would fall into what Greller and Drachsler characterize as "potential dangers" of educational data [23].

Previous empirical evidence suggests that students are heavily influenced by grade-prediction tools when making enrollment choices [41]. That is, students experience a sort of cognitive commitment by which they trust, almost blindly, what the tool predicts about their future performance. We did not observe a similar effect on the advisors who, conversely, tend to rely on their own experience and personal views of the courses. Besides, advisors have an innate tendency to account for the students' extracurricular circumstances, which is vital for proper academic advising. It follows that appropriate human judgment should not be replaced by a grade prediction tool, because advising students implies integrating non-quantitative information that cannot be handled by a tool like iCoRA. The academic advisors' expertise is therefore paramount to address the potential negative impacts of grade-based AI tools while exploiting its virtues for the sake of a more effective student-advisor dialogue.

7 LIMITATIONS AND FUTURE WORK

Our findings are limited by the size and characteristics of our sample. We deliberately set out to test participants from a specific CS undergraduate program. While small ($n = 15$), our sample included 52% of the departmental advisors of our research site. We acknowledge, however, that this group may not be representative of all potential audiences. In particular, our participants were familiar

with prediction models, something that cannot be taken for granted for advisors from other study programs. Further empirical work is still needed to fully understand the influence of grade predictions on academic advisors from other areas. These investigations could take a comparative approach to uncover differences in the advising processes with and without a tool like iCoRA.

Another limitation of our experimental design is that advisors made course recommendations in the absence of students, something that does not happen in the actual advising sessions conducted at our research site. This meant that our participants had access to the students' academic and enrollment history, without any knowledge of the students' personal circumstances. We decided in favor of this design to study the effects of iCoRA in isolation, without the confounding factor that the students' unique circumstances may have introduced. Along the same lines, in our experiment, we used synthetic student profiles rather than data from real students. We did so to have full control of the experimental performance profiles and to make sure these were clearly distinguishable. This also allowed us to confront advisors to students with the same enrollment options. This study design choice, however, may have an impact on the generalizability and validity of our results, because, in reality, students may exhibit phases of low/high performance, or may have more affinity to specific areas of their study program. It could also be argued that the marked performance profiles we used may have influenced how much advisors focused on the students' GPA during the recommendation tasks. Notwithstanding, we highlight that during the interviews, our participants described the GPA-driven advising strategy as their usual approach to course recommendation. This is not surprising because the institutional policies of our research site rely heavily on the GPA as an academic performance metric. This is illustrated, for example, by the academic advising policies described in Section 3.1. Despite the prominence of the GPA in the current institutional policies of our research site, our participants seemed fully aware of the limitations of this metric as a proxy for academic success, so much so that they deemed their meetings with students as an opportunity to mitigate the GPA's potential negative effects. In this regard, our research work has helped to raise voices for the use of alternative, more comprehensive—and less troublesome—metrics of student performance (e.g., mastery learning approaches [9, 25], portfolios [39, 55]), or simply assessing performance via ungraded tasks [33, 45]. The long-term effects and implications of these changes, however, deserve their own space and are outside the scope of this paper.

Finally, while our observations suggest that a tool like iCoRA may require the introduction of mechanisms to build trust, our study does not provide answers on how to design or implement those mechanisms. Further research efforts are still needed to investigate the requirements for AI-based tools to be successfully adopted in the context of academic advising and course recommendation.

8 CONCLUSION

This paper investigated the effects that grade predictions have on academic advisors during term planning. To this end, academic advisors of a CS undergraduate program within an Engineering-oriented HEI used a tool that supports the composition of course sets and provides grade predictions and explanations thereof. The

advisors used the tool to recommend courses for the forthcoming academic term of three different students. The students had the same set of courses available for enrollment, but they exhibited distinguishable academic performance profiles, namely low, average, and high GPA.

Our observations show that the advisors' self-reported usual approach to course recommendation is barely influenced by the presence of the predicted grades. Instead, advisors mostly rely on their experience and personal views about the courses' workload, time commitments, and relative importance. So much so that some advisors even disregarded the predictive support of the tool altogether. We also found that advisors tend to make more challenging recommendations to high-performing students, even in the presence of pessimistic predictions. Our results also suggest that advisors may spend more time interacting with the grade prediction tool when recommending courses to low-performing advisees, attempting to maximize the potential GPA gain for this kind of student.

These observations highlight the importance of advisors in any academic support system, particularly given previous empirical evidence that shows that students are heavily influenced by grade predictions during course selection. Although advisors valued the predictive support to make more informed recommendations, they also raised concerns about the trustworthiness and validation of the predicted grades. These insights may guide the design and development of new data-based supporting tools for academic advising.

ACKNOWLEDGMENTS

The authors wish to thank the reviewers and associate chairs for their advice and insightful comments. We would also like to express our gratitude to the academic advisors who generously donated their time to participate in this study. This research was supported and partially financed by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No. 952215.

REFERENCES

- [1] The Quality Assurance Agency. 2001. *QAA Code of Practice for the assurance of academic quality and standards in higher education: Career Education, Information and Guidance (CEIG)*. Taylor & Francis. <https://doi.org/10.11120/plan.2001.00020026>
- [2] Abdulrahman Alkhoori, Mohammad Amin Kuhail, and Abdulla Alkhoori. 2020. UniBud: A Virtual Academic Adviser. In *2020 12th Annual Undergraduate Research Conference on Applied Computing (URC)*. 1–4. <https://doi.org/10.1109/URC49805.2020.9099191>
- [3] Georgina Argüello and María Grethel Méndez. 2019. Virtual advising: A tool for retention, engagement, and success for the graduate distance learner. *Distance Learning* 16, 2 (2019), 51–57. <https://eric.ed.gov/?id=EJ1299512>
- [4] Michael Backenköhler, Felix Scherzinger, Adish Singla, and Verena Wolf. 2018. Data-Driven Approach towards a Personalized Curriculum. *International Educational Data Mining Society* (2018). <https://eric.ed.gov/?id=ED593214>
- [5] Farhad Bilal Baha'Addin. 2013. To Meet The Academic Advising Needs Of The Students In A More Interactive And Effective Way. *International Journal of Engineering Research & Technology* 2, 4 (2013). <https://doi.org/10.17577/IJERTV2IS4715>
- [6] Bahr. 2008. Cooling Out in the Community College: What is the Effect of Academic Advising on Students' Chances of Success? *Research in Higher Education* 49, 8 (Jul 2008), 704–732. <https://doi.org/10.1007/s11162-008-9100-0>
- [7] Talia Bar, Vrinda Kadiyali, and Asaf Zussman. 2008. Quest for Knowledge and Pursuit of Grades: Information, Course Selection, and Grade Inflation. *Behavioral & Experimental Economics* (2008). <https://doi.org/10.2139/ssrn.1019580>
- [8] Talia Bar, Vrinda Kadiyali, and Asaf Zussman. 2009. Grade Information and Grade Inflation: The Cornell Experiment. *Journal of Economic Perspectives* 23, 3 (September 2009), 93–108. <https://doi.org/10.1257/jep.23.3.93>
- [9] James H Block and Robert B Burns. 1976. Mastery learning. *Review of research in education* 4 (1976), 3–49.
- [10] Robert Bodily and Katrien Verbert. 2017. Review of Research on Student-Facing Learning Analytics Dashboards and Educational Recommender Systems. *IEEE Transactions on Learning Technologies* 10, 4 (2017), 405–418. <https://doi.org/10.1109/TLT.2017.2740172>
- [11] Richard E Boyatzis. 1998. *Transforming Qualitative Information: Thematic Analysis and Code Development*. Sage.
- [12] Hana Bydžovská. 2016. Course Enrollment Recommender System. *International Educational Data Mining Society* (2016). <https://eric.ed.gov/?id=ED592681>
- [13] Jaime Castells, Mohammad Poul Doust, Luis Galárraga, Gonzalo Gabriel Méndez, Margarita Ortiz-Rojas, and Alberto Jiménez. 2020. A Student-oriented Tool to Support Course Selection in Academic Counseling Sessions. In *Proceedings of the Workshop on Adoption, Adaptation and Pilots of Learning Analytics in Underrepresented Regions, co-located with the 15th European Conference on Technology Enhanced Learning 2020 (ECTEL 2020)*. <http://ceur-ws.org/Vol-2704/paper4.pdf>
- [14] Zenobia C.Y. Chan, Ho Yan Chan, Hang Chak Jason Chow, Sze Nga Choy, Ka Yan Ng, Koon Yiu Wong, and Pui Kan Yu. 2019. Academic advising in undergraduate education: A systematic review. *Nurse Education Today* 75 (2019), 58–74. <https://doi.org/10.1016/j.nedt.2019.01.009>
- [15] Sven Charleer, Andrew Vande Moere, Joris Klerkx, Katrien Verbert, and Tinne De Laet. 2018. Learning Analytics Dashboards to Support Adviser-Student Dialogue. *IEEE Transactions on Learning Technologies* 11, 3 (2018), 389–399. <https://doi.org/10.1109/TLT.2017.2720670>
- [16] Sorathan Chaturapruek, Thomas S. Dee, Ramesh Johari, René F. Kizilcec, and Mitchell L. Stevens. 2018. How a Data-Driven Course Planning Tool Affects College Students' GPA: Evidence from Two Field Experiments. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale (London, United Kingdom) (L@S '18)*. Association for Computing Machinery, New York, NY, USA, Article 63, 10 pages. <https://doi.org/10.1145/3231644.3231668>
- [17] Leah Chong, Guanglu Zhang, Kosa Goucher-Lambert, Kenneth Kotovsky, and Jonathan Cagan. 2022. Human confidence in artificial intelligence and in themselves: The evolution and impact of confidence on adoption of AI advice. *Computers in Human Behavior* 127 (2022), 107018. <https://doi.org/10.1016/j.chb.2021.107018>
- [18] European Commission, Sport Directorate-General for Education, Youth, Culture, S Wollscheid, B Stensaker, B Jongbloed, H Vossensteyn, L Cremonini, E Hovdhagen, F Kaiser, and A Kottmann. 2015. *Dropout and completion in higher education in Europe : main report*. Publications Office. <https://doi.org/doi/10.2766/826962>
- [19] Joe Cuseo. 2003. Academic advisement and student retention: Empirical connections and systemic interventions. *National Academic Advising Association* (2003), 2019–01.
- [20] Jayne K. Drake. 2011. The Role of Academic Advising in Student Retention and Persistence. *About Campus* 16, 3 (2011), 8–12. <https://doi.org/10.1002/abc.20062>
- [21] Steven C. Ender, Roger B. Winston Jr., and Theodore K. Miller. 1982. Academic advising as student development. *New Directions for Student Services* 1982, 17 (1982), 3–18. <https://doi.org/10.1002/ss.37119821703> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/ss.37119821703
- [22] Omri Gillath, Ting Ai, Michael S. Branicky, Shawn Keshmiri, Robert B. Davison, and Ryan Spaulding. 2021. Attachment and trust in artificial intelligence. *Computers in Human Behavior* 115 (2021), 106607. <https://doi.org/10.1016/j.chb.2020.106607>
- [23] Wolfgang Greller and Hendrik Drachsler. 2012. Translating Learning into Numbers: A Generic Framework for Learning Analytics. *Journal of Educational Technology & Society* 15, 3 (2012), 42–57. <http://www.jstor.org/stable/jeductechsoci.15.3.42>
- [24] Yu Guo, Yue Chen, Yuanyan Xie, and Xiaojuan Ban. 2022. An Effective Student Grouping and Course Recommendation Strategy Based on Big Data in Education. *Information* 13, 4 (2022). <https://doi.org/10.3390/info13040197>
- [25] Thomas R Guskey. 2022. *Implementing mastery learning*. Corwin Press.
- [26] Francisco Gutiérrez, Karsten Seipp, Xavier Ochoa, Katherine Chiluiza, Tinne De Laet, and Katrien Verbert. 2018. LADA: A Learning Analytics Dashboard for Academic Advising. *Computers in Human Behavior* 107 (2018). <https://doi.org/10.1016/j.chb.2018.12.004>
- [27] Wesley R Habley. 2004. *The status of academic advising: Findings from the ACT sixth national survey*. National Academic Advising Association. <https://nacada.ksu.edu/Resources/Journal/Current-Past-Book-Reviews/The-Status-of-Academic-Advising-Findings-from-the-ACT-Sixth-National-Survey.aspx>
- [28] Heba Ismail, Nada Hussein, Rawan Elabyad, and Salma Said. 2021. A Serverless Academic Adviser Chatbot. In *The 7th Annual International Conference on Arab Women in Computing in Conjunction with the 2nd Forum of Women in Research (Sharjah, United Arab Emirates) (ArabWIC 2021)*. Association for Computing Machinery, New York, NY, USA, Article 27, 5 pages. <https://doi.org/10.1145/3485557.3485587>
- [29] Sandeep M. Jayaprakash, Erik W. Moody, Eitel J.M. Lauria, James R. Regan, and Joshua D. Baron. 2014. Early Alert of Academically At-Risk Students: An Open Source Analytics Initiative. *Journal of Learning Analytics* 1, 1 (May 2014), 6–47. <https://doi.org/10.18608/jla.2014.11.3>
- [30] Weijie Jiang and Zachary A. Pardos. 2021. Towards Equity and Algorithmic Fairness in Student Grade Prediction. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (Virtual Event, USA) (AIES '21)*. Association

- for Computing Machinery, New York, NY, USA, 608–617. <https://doi.org/10.1145/3461702.3462623>
- [31] Weijie Jiang, Zachary A. Pardos, and Qiang Wei. 2019. Goal-Based Course Recommendation. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge* (Tempe, AZ, USA) (LAK19). Association for Computing Machinery, New York, NY, USA, 36–45. <https://doi.org/10.1145/3303772.3303814>
 - [32] Ahmad A. Kardan, Hamid Sadeghi, Saeed Shiry Ghidary, and Mohammad Reza Fani Sani. 2013. Prediction of student course selection in online higher education institutes using neural network. *Computers and Education* 65 (2013), 1–11. <https://doi.org/10.1016/j.compedu.2013.01.015>
 - [33] Alfie Kohn and Susan D Blum. 2020. *Ungrading: Why Rating Students Undermines Learning (and What to Do Instead)*. West Virginia University Press.
 - [34] Mohammad Amin Kuhail, Haseena Al Katheeri, Joao Negreiros, Ahmed Seffah, and Omar Alfandi. 2022. Engaging Students With a Chatbot-Based Academic Advising System. *International Journal of Human-Computer Interaction* 0, 0 (2022), 1–27. <https://doi.org/10.1080/10447318.2022.2074645>
 - [35] Markus Langer and Richard N. Landers. 2021. The future of artificial intelligence at work: A review on effects of decision automation and augmentation on workers targeted by algorithms and third-party observers. *Computers in Human Behavior* 123 (2021), 106878. <https://doi.org/10.1016/j.chb.2021.106878>
 - [36] Mei Shyan Lim, Sin-Ban Ho, and Ian Chai. 2021. Design and Functionality of a University Academic Advisor Chatbot as an Early Intervention to Improve Students' Academic Performance. In *Computational Science and Technology*, Rayner Alfred, Hiroyuki Iida, Havaluddin Havaluddin, and Patricia Anthony (Eds.). Springer Singapore, Singapore, 167–178. https://doi.org/10.1007/978-981-33-4069-5_15
 - [37] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 4768–4777. <https://doi.org/10.5555/3295222.3295230>
 - [38] N D Lynn and A W R Emanuel. 2021. A review on Recommender Systems for course selection in higher education. *IOP Conference Series: Materials Science and Engineering* 1098, 3 (mar 2021), 032039. <https://doi.org/10.1088/1757-899x/1098/3/032039>
 - [39] Effie Maclellan. 2004. How convincing is alternative assessment for use in higher education? *Assessment & Evaluation in Higher Education* 29, 3 (2004), 311–321. <https://doi.org/10.1080/0260293042000188267>
 - [40] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.* 54, 6, Article 115 (jul 2021), 35 pages. <https://doi.org/10.1145/3457607>
 - [41] Gonzalo Gabriel Méndez, Luis Galárraga, and Katherine Chiluiza. 2021. Showing Academic Performance Predictions during Term Planning: Effects on Students' Decisions, Behaviors, and Preferences. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 22, 17 pages. <https://doi.org/10.1145/3411764.3445718>
 - [42] Quanza E. Mooring. 2016. Recruitment, advising, and retention programs — Challenges and solutions to the international problem of poor nursing student retention: A narrative literature review. *Nurse Education Today* 40 (2016), 204–208. <https://doi.org/10.1016/j.nedt.2016.03.003>
 - [43] Mirna Nachouki and Mahmoud Abou Naaj. 2019. Process Automation Tool for Academic Advising. In *2019 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT)*. 1–6. <https://doi.org/10.1109/ISSPIT47144.2019.9001864>
 - [44] Jaideep T Naidu and Chae Mi Lim. 2020. Optimizing Grade Point Averages During the Pandemic at a Regional University. *Business Education Innovation Journal* 12, 2 (December 2020), 7.
 - [45] Jennifer R Newton, Mira Cole Williams, and Danielle M Feeney. 2020. Implementing non-traditional assessment strategies in teacher preparation: Opportunities and challenges. *Journal of Culture and Values in Education* 3, 1 (2020), 39–51.
 - [46] Viet Anh Nguyen, Hoa-Huy Nguyen, Duc-Loc Nguyen, and Minh-Duc Le. 2021. A course recommendation model for students based on learning outcome. *Education and Information Technologies* 26, 5 (2021), 5389–5415. <https://doi.org/10.1007/s10639-021-10524-0>
 - [47] Deborah M. Oh, Simeon Slovacsek, Susan Tucker, and Ann Hafner. 2003. Assessment Outcomes of Pre-service Teachers. *Assessment & Evaluation in Higher Education* 28, 3 (2003), 279–295. <https://doi.org/10.1080/0260293032000059630>
 - [48] Emmanuel Okewu and Olawande Daramola. 2017. Design of a learning analytics system for academic advising in Nigerian universities. In *2017 International Conference on Computing Networking and Informatics (ICCNi)*. 1–8. <https://doi.org/10.1109/ICCNi.2017.8123785>
 - [49] Idowu Dauda Oladipo, Joseph Bamidele Awotunde, Muyideen AbdulRaheem, Oluwasegun Osemudiamé Ige, Ghaniyyat Bolanle Balogun, Adekola Rasheed Tomori, and Fatimoh Abidemi Taofeek-Ibrahim. 2021. An Improved Course Recommendation System Based on Historical Grade Data Using Logistic Regression. In *Applied Informatics*, Hector Florez and Ma Florencia Pollo-Cattaneo (Eds.). Springer International Publishing, Cham, 207–221. https://doi.org/10.1007/978-3-030-89654-6_15
 - [50] Pardos, Fan, and Jiang. 2019. Connectionist recommendation in the wild: on the utility and scrutability of neural networks for personalized course guidance. *User Modeling and User-Adapted Interaction* 29, 2 (Feb 2019), 487–525. <https://doi.org/10.1007/s11257-019-09218-7>
 - [51] Zachary A. Pardos and Weijie Jiang. 2020. Designing for Serendipity in a University Course Recommendation System. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge* (Frankfurt, Germany) (LAK '20). Association for Computing Machinery, New York, NY, USA, 350–359. <https://doi.org/10.1145/3375462.3375524>
 - [52] Elizabeth D. Phillips. 2013. Improving Advising Using Technology and Data Analytics. *Change: The Magazine of Higher Learning* 45, 1 (2013), 48–55. <https://doi.org/10.1080/00091383.2013.749151>
 - [53] Richard W Schwartz, James E Burgett, Amy V Blue, Michael B Donnelly, and David A Sloan. 1997. Problem-based learning and performance-based testing: Effective alternatives for undergraduate surgical education and assessment of student performance. *Medical Teacher* 19, 1 (1997), 19–23. <https://doi.org/10.3109/01421599709019341>
 - [54] Bruce S. Sharkin. 2004. College Counseling and Student Retention: Research Findings and Implications for Counseling Centers. *Journal of College Counseling* 7, 2 (2004), 99–108. <https://doi.org/10.1002/j.2161-1882.2004.tb00241.x> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.2161-1882.2004.tb00241.x
 - [55] Katrien Struyven, Filip Dochy, Steven Janssens, Wouter Schelfhou, and Sarah Gielens. 2006. The overall effects of end-of-course assessment on student performance: A comparison between multiple choice testing, peer assessment, case-based assessment and portfolio assessment. *Studies in Educational Evaluation* 32, 3 (2006), 202–222. <https://doi.org/10.1016/j.stueduc.2006.08.002>
 - [56] Lemaro R. Thompson and Leon C. Prieto. 2013. Improving Retention Among College Students: Investigating the Utilization of Virtualized Advising. *Academy of Educational Leadership Journal* 17, 4 (2013), 13–26. <https://www.proquest.com/scholarly-journals/improving-retention-among-college-students/docview/1462525754/se-2> Copyright - Copyright Jordan Whitney Enterprises, Inc 2013; Document feature - Tables; Diagrams; ; Last updated - 2021-09-09; SubjectsTermNotLitGenreText - United States-US.
 - [57] Chris van Klaveren, Karen Kooiman, Ilja Cornelisz, and Martijn Meeter. 2019. The Higher Education Enrollment Decision: Feedback on Expected Study Success and Updating Behavior. *Journal of Research on Educational Effectiveness* 12, 1 (2019), 67–89. <https://doi.org/10.1080/19345747.2018.1496501>
 - [58] Camilo Vieira, Paul Parsons, and Vetria Byrd. 2018. Visual Learning Analytics of Educational Data: A Systematic Literature Review and Research Agenda. *Computers & Education* 122 (2018), 119–135. <https://doi.org/10.1016/j.compedu.2018.03.018>
 - [59] Graham Whitehead. 2013. Developing institutional strategies to support failing/failed part-time students in higher education. *The Journal of Practice Teaching and Learning* 11, 2 (Mar. 2013), 27–46. <https://doi.org/10.1921/jpts.v11i2.265>
 - [60] Jamie L. Workman. 2015. Exploratory Students' Experiences With First-Year Academic Advising. *NACADA Journal* 35, 1 (07 2015), 5–12. <https://doi.org/10.12930/NACADA-14-005>
 - [61] Adena D. Young-Jones, Tracie D. Burt, Stephanie Dixon, and Melissa J. Hawthorne. 2013. Academic advising: does it really impact student success? *Quality Assurance in Education* 21, 1 (Jan 2013), 7–19. <https://doi.org/10.1108/09684881311293034>
 - [62] Run Yu, Zach Pardos, Hung Chau, and Peter Brusilovsky. 2021. Orienting Students to Course Recommendations Using Three Types of Explanation. In *Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization* (Utrecht, Netherlands) (UMAP '21). Association for Computing Machinery, New York, NY, USA, 238–245. <https://doi.org/10.1145/3450614.3464483>
 - [63] Qing Zhou and Fang Yu. 2008. Knowledge-Based Major Choosing Decision Making for Remote Students. In *International Conference on Computer Science and Software Engineering*, Vol. 5. 474–478. <https://doi.org/10.1109/CSSE.2008.379>