Course Recommender Systems with Statistical Confidence

Zachary Warnes
Department of Data Science
and Knowledge Engineering,
Maastricht University, Maastricht 6200MD,
The Netherlands
z.warnes@student.maastrichtuniversity.nl

Evgueni Smirnov
Department of Data Science
and Knowledge Engineering,
Maastricht University, Maastricht 6200MD,
The Netherlands
smirnov@maastrichtuniversity.nl

ABSTRACT
Selecting courses in an open-curriculum education program is a difficult task for students and academic advisors. Course recommendation systems nowadays can be used to reduce the complexity of this task. To control the recommendation error, we argue that course recommendations need to be provided together with *statistical* confidence. The latter can be used for computing a statistically valid set of recommended courses that contains courses a student is likely to take with a probability of at least $1-\epsilon$ for a user-specified significance level ϵ. For that purpose, we introduce a generic algorithm for course recommendation based on the conformal prediction framework. The algorithm is used for implementing two conformal course recommender systems. Through experimentation, we show that these systems accurately suggest courses to students while maintaining statistically valid sets of courses recommended.

Keywords
Recommender Systems, Course Recommendation, Conformal Prediction

1. INTRODUCTION
Recommender systems are systems capable of predicting the preferences of users over sets of items [1]. They can be found almost everywhere in the digital space, shaping the choices we make, the products we buy, the books we read, or the movies we watch. The range of applications of recommender systems has been broadened recently to the education domain, especially in higher education [5]. There are systems reported that provide recommendations for academic choices, learning activities, learning resources, and learning collaborations [14].

Among the recommender systems for academic choices, there exists a particular interest in systems that recommend courses [3]. There is a wide range of such systems that differ in the underlying recommendation mechanism, accuracy, type of recommendations (courses, course sequences, course concentrations), and type of representation. It has been recently recognized that course recommender systems need to be safe [11]; i.e., course recommendations need to be provided with confidence information that will help a student to make a better course selection. There exist different approaches to delivering such confidence information from course preference ranks estimated by the underlying recommendation mechanisms [3, 6, 10, 12] to separate warning modules [11]. The characteristic feature of these approaches is that they are heuristic, and thus they do not provide any theoretical guarantees for the quality of course recommendation.

In this paper, we argue that course recommendations need to be supported with *statistical* confidence. This confidence will allow computing a statistically valid set of recommended courses that contains courses a student is likely to take with a probability of at least $1 - \epsilon$ for a user-specified significance level ϵ. To achieve this, we employ the well-known conformal-prediction framework [4, 15, 16]. We design a generic algorithm for conformal course recommendation capable of computing statistically valid sets of courses for students. The algorithm is used for implementing two conformal course recommender systems that employ a content-based recommendation mechanism. The first system is instance-based, and the second system is an exemplar-based system [13].

The conformal course recommender systems have been implemented for academic advising of University College Maastricht, a Liberal Arts Bachelor study with an open curriculum. In this study, students personalize their program by selecting courses that align with their academic and personal interests. In total, students choose around 40 out of the 160 possible educational modules; i.e., they create a program by selecting one path out $\binom{160}{40}$ possible. Our recommender systems are tested to facilitate this process. The initial experimental results show that the systems accurately recommend courses while providing statistically valid sets of courses recommended.

The rest of the paper is organized as follows. The related work is provided in Section 2. Section 3 formalizes the task of course recommendation. The course and student topic models used for course recommendation are briefly described in Section 4. Section 5 introduces the generic algorithm for conformal course recommendation and its instantiations: the instance-based and exemplar-based recommender sys-

Zachary Warnes and Evgueni Smirnov "Course Recommender Systems with Statistical Confidence" In: *Proceedings of The 13th International Conference on Educational Data Mining (EDM 2020)*, Anna N. Rafferty, Jacob Whitehill, Violetta Cavalli-Sforza, and Cristobal Romero (eds.) 2020, pp. 509 - 515

tems. The student-course data is described in Section 6. Section 7 provides the experiments and discussion. Finally, Section 8 concludes the paper.

2. RELATED WORK

Course recommender systems received significant attention since the very first publications [12, 18, 17]. Meanwhile, these systems have become very diverse. Following the main trends in recommender-system research there are different types of course recommender systems: content-based systems [10, 11], collaborative-filtering systems [3], hybrid systems [3, 6], and popularity-based ranking systems [6]. Most of these systems are capable of providing (explicitly/implicitly) confidence information for course recommendations. However, this does not give any guarantee for the quality of course recommendation in a statistical sense.

Confidence-based recommender systems have been proposed based on collaborative filtering. The first system is based on group recommender systems [8], and the second one is based on matrix factorization [7]. Both systems can be directly applied for course recommendations, however, under assumptions typical for collaborative filtering. For example, plenty of data is available; the course order does not matter. In this context, we note that we propose a generic algorithm for conformal course recommendation that is not tailored to the recommendation mechanism: collaborative filtering or content-based filtering. The only requirement to apply this algorithm is to have a function that estimates the typicality of a course w.r.t. other courses taken by a student (see conformity functions in Section 5).

3. RECOMMENDATION TASK WITH CONFIDENCE

Let T be a set of topics t considered in a set C courses c. To indicate the degree of presence of topic t in course $c \in C$ we employ weight $w_{c,t}$. Topic weights $w_{c,t}$ of course $c \in C$ represents a topic model of this course. We assume that the topic models (i.e. topics' weights $w_{c,t}$) are provided initially for all the courses $c \in C$. We describe our approach to derive these models in the next Section.

The courses $c \in C$ are given for a set S of students s. To indicate the degree student $s \in S$ masters topic t we employ weight $w_{s,t}$. Topic weights $w_{s,t}$ of student $s \in S$ represents a topic model of the student w.r.t. courses $c \in C$. Thus, they are computed w.r.t. set C_s of courses student s has taken; i.e. for any topic $t \in T$ we have:

$$w_{s,t} = \frac{\sum_{c \in Cs} w_{c,t}}{|Cs|}, \tag{1}$$

If we assume a specific ordering of the topics $t \in T$, then:

- the topics' weights $w_{c,t}$ for course c form a topic-model vector w_c for c, and
- the topics' weights $w_{s,t}$ for student s form a topic-model vector w_s for s.

The topic-model vectors w_c and w_s "live" in the same space W. Due to the number $|T|$ of all the topics, we employ the cosine similarity over W. It can be used to compute similarity for any two topic-model vectors that represent courses and students.

The topic-model vectors w_c of all the courses $c \in C$ form the course data set W_C defined as $\{w_c\}_{c \in C}$. Analogously, the topic-model vector w_t of all the students $t \in T$ form the student data set W_S defined as $\{w_t\}_{t \in T}$. In this context, we introduce the recommendation task considered in this paper. Given a course data set W_C, a student data set W_S, and a student $s \in S$ with topic-model vector $w_s \in W_S$, the task is to compute a recommendation set $C_s^\epsilon \subseteq C \setminus C_s$ that contains courses that indeed fit student s with a probability at least $1 - \epsilon$ for a predefined significance level $\epsilon \in [0, 1]$.

4. COURSE AND STUDENT TOPIC MODELLING

We employed the topic-modeling approach proposed in [11]. The set T of topics t was identified from the course descriptions using the Latent Dirichlet Allocation (LDA) generative model [2]. Each topic $t \in T$ is given by a probability distribution over the vocabulary derived from all the descriptions. Thus, each course $c \in C$ is represented by topics t, which words are present in the description of that course. Student topic models are derived based on the topics courses using formula (1).

5. CONFORMAL COURSE RECOMMENDATION

This section introduces a conformal course recommendation. First, we present a conformal test for course inclusion and a generic algorithm for conformal course recommender systems. Then we provide two instantiations of this algorithm.

5.1 Generic Conformal Course Recommender

Consider a particular student $s \in S$ with her set C_s of courses. We assume that student s is represented by a probability distribution P_s; i.e. P_s has generated the course set C_s for s. Thus, to decide whether to recommend a new course $c \notin C_s$ for student s, we perform a statistical test of the null hypothesis that the set $C_s \cup \{c\}$ is generated by the student distribution P_s under the exchangeability assumption [15] [1].

We implement the statistical test according to the conformal-prediction framework [15]. It makes use of course conformity scores. The conformity score α_c of a course c is defined as a score that indicates how typical c in set $C_s \cup \{c\}$. The conformity score α_c is computed by a course conformity function A. The latter is a mapping from $2^C \times C$ to $\mathbb{R} \cup \{+\infty\}$; i.e. it returns for any course set C_s and any course c a score α_c that indicates how typical is course c for the courses in $C_s \cup \{c\}$. Depending on the implementation of the conformity function for the course and student topic models, we can have recommender systems based on content/collaborative filtering (See the next section).

The conformity score α_c of a new course c is used as a test statistic for the null hypothesis that the set $C_s \cup \{c\}$ is generated by the student distribution P_s according to the

[1]We note that the exchangeability assumption is weaker than the well-know i.i.d. assumption.

exchangeability assumption. The p-value p_c for the null hypothesis, is calculated as the fraction of the courses in $C_s \cup \{c\}$ associated with conformity scores that are equal to or smaller than α_c. The larger the value of p_c, the more likely it is to observe the value of α_c under the null hypothesis, and the more confidence we have in course c. If we set a course significance level ϵ_c (probability of the error) in range of $[0, 1]$, then the statistical test will accept the null hypothesis if $p_c > \epsilon_c$ and course c will be recommended.

If we perform the statistical test from above for student $s \in S$ over all n courses from set $C \setminus C_s$ (that s has not taken), then we can compute a recommended set of courses. We summarize this process in Algorithm 1 given below. It presents a generic conformal course recommender algorithm. Given a course significance level $\epsilon_c \in [0, 1]$ and set C_s of m courses taken by student s, the algorithm computes set $C_s^\epsilon \subseteq C \setminus C_s$ of recommended courses for student s on the chosen significance level ϵ. To decide whether to include course $c_i \in C \setminus C_s$ in set C_s^ϵ the algorithm computes the conformity score α_{c_j} for each course $c_j \in C_s \cup \{c_i\}$ using the nonconformity function A (step 4). The conformity scores α_{c_j} are used for computing the p-value p_{c_i} of course c_i (step 6). Once p_{c_i} has been obtained, course c_i is added to the set C_s^ϵ of recommended courses if $p_{c_i} > \epsilon_c$ (step 7). This process is repeated for every course $c_i \in C \setminus C_s$.

We note that the generic conformal course recommender algorithm computes valid recommendation sets C_s^ϵ for a significance level ϵ that usually is bigger than the course significance level ϵ_c (used in Algorithm 1). To explain this phenomena we follow the approach proposed in [9]. W.l.o.g. assume that the set $C \setminus C_s$ is the same for all the students $s \in S$. Let e_{c_i} be a random error variable for a course c_i from the n courses in $C \setminus C_s$. The variable e_{c_i} equals 1 if course c_i does not fit student s; and 0, otherwise. Assume that we set the course significance level ϵ_c so that $p(e_{c_1} = 1) < \epsilon_c, p(e_{c_2} = 1) < \epsilon_c, \ldots, p(e_{c_n} = 1) < \epsilon_c$. This implies that the expected number of courses incorrectly recommended, $e_{c_1} + e_{c_2} + \ldots + e_{c_n}$, is bounded by $n\epsilon_c$; i.e. $\mathbf{E}[\sum_{c_i \in C \setminus C_s} e_{c_i}] \leq n\epsilon_c$. If we know number t of courses from $C \setminus C_s$ that fit student s in advance, then:

$$\frac{1}{t}\mathbf{E}[\sum_{c_i \in C \setminus C_s} e_{c_i}] \leq \frac{n}{t}\epsilon_c. \tag{2}$$

We note that $\frac{1}{t}\mathbf{E}[\sum_{c_i \in C \setminus C_s} e_{c_i}]$ is the expected error and $\frac{n}{t}\epsilon_c$ is a significance level ϵ for which validity of recommendation sets C_s^ϵ can be established. This implies:

$$\epsilon_c = \frac{t}{n}\epsilon \tag{3}$$

Thus, to guarantee valid recommendation sets $C_s^\epsilon \subseteq C \setminus C_s$ that contains courses that fit students with a probability at least $1 - \epsilon$ we need to set the course significance level ϵ_c according to formula (3) when we initialize the generic conformal course recommender algorithm from Algorithm 1.

Algorithm 1 Generic Conformal Course Recommender

Input: Course significance level ϵ_c,
Set C_s of m courses taken by student s.
Output: Set C_s^ϵ of recommended courses for student s.

1: Set course set C_s^ϵ equal to $\emptyset$.
2: **for each** course $c_i \in C \setminus C_s$ **do**
3: **for each** course $c_j \in C_s \cup \{c_i\}$ **do**
4: Set conformity score α_{c_j} of course c_j equal to $A(C_s \cup \{c_i\}, c_j)$.
5: **end for**
6: Set p_{c_i} equal to $\frac{\#\{c_j \in C_s \cup \{c_i\} | \alpha_{c_j} \leq \alpha_{c_i}\}}{m+1}$.
7: Add course c_i to C_s^ϵ if $p_{c_i} > \epsilon_c$.
8: **end for**
9: Output set $C_s^{\epsilon_c}$ of recommended courses for student s.

To establish the validity of sets C_s^ϵ of recommended courses, we adapt the error metric from [9, 19]. Assume that for student $s \in S$, we have a test set of courses, and we know that within this set, there is a true set C_s^t of courses that student s will take. We define the individual error e_s for student $s \in S$ as the proportion of the courses in the true set C_s^t that are not recommended, i.e.

$$e_s = \frac{|C_s^t \setminus (C_s^\epsilon \cap C_s^t)|}{|C_s^t|}.$$

In this context, the error e of a conformal course recommender is defined as the averaged error e over all the students $s \in S$:

$$e = \frac{\sum_{s \in S} e_s}{|S|}.$$

We note that the individual error e_s corresponds to the expected error $\frac{1}{|C_s^t|}\mathbf{E}[\sum_{c_i \in C \setminus C_s} e_{c_i}]$. Thus, to show experimentally that a conformal course recommender is valid for any significance level ϵ in $[0, 1]$, we have to show that the error e is less than or equal to ϵ.

The validity of a conformal course recommender can be trivially achieved if the recommender outputs all the possible courses from set $C \setminus C_s$. Thus, we need to estimate the informational efficiency of the recommender. For this purpose we employ the size SR of the recommended set C_s^ϵ of courses averaged over all the students:

$$SR = \frac{\sum_{s \in S} |C_s^\epsilon|}{|S|}.$$

5.2 Content-based Conformal Course Recommender Systems

The generic conformal course recommender algorithm can be instantiated if we specify the course conformity function A. This function can be done using different recommender mechanisms, e.g., collaborative filtering or content-based filtering. In this paper, we assume the existence of topic model vectors of courses and students that fit the content-based filtering scenario (see Section 3). That is why we propose conformity functions for two content-based conformal course recommender systems specified below.

The first system is an instance-based conformal course recommender system (ICCRS). Any student $s \in S$ is represented by a set of topic-model vectors (instances) w_c of the courses $c \in C_s$ she has taken. In this context the course conformity function A outputs for any course $c \in C$ and course set C_s of student $s \in S$ an averaged similarity of c with courses in C; i.e. $\frac{1}{|C_s \setminus \{c\}|} \sum_{c' \in C_s \setminus \{c\}} \cos(w_c, w_{c'})$ where cos is the cosine similarity.

The second system is an exemplar-based conformal course recommender system (ECCRS). It employs topic-model vector w_s (exemplar) of student $s \in S$ computed using formula (1). In this context the course conformity function A outputs for any course $c \in C$ and course set C_s of student $s \in S$ a value equal to $\cos(w_c, w_s)$, where topic-model vector w_s of student $s \in S$ is based on the courses in $C_s \setminus \{c\}$ and cos is the cosine similarity.

The computational complexity of ECCRS is higher than that of ICCRS since, for any student, we need to recompute her topic-model vectors w_s by excluding courses one by one. However, ECCRS has better explanation capabilities. The topic-model vector w_s of student s represents the current levels of topic mastering, and the topic-model vector w_c, of course, c represents the topics covered in the course. Thus, the cosine match can explain why the course has been selected/rejected.

6. STUDENT-COURSE DATA

ICCRS and ECCRS have been implemented as course recommender systems for University College Maastricht. The college has provided course enrollment data from 2008 to 2017. This data includes course and student identifiers, grades for each course, details regarding course assessment, ECT credits, and course descriptions. The course descriptions facilitate the construction of topic values for both the student model and the course model. The calculation of topic values is with LDA, and an optimal number of topics is determined through maximum likelihood estimation. This optimization results in sixty-five topic areas representing the course catalog [11]. We remove modules without descriptions from consideration. In total, 143 courses and 2422 students enrolled in at least one course remain.

The rates of course enrollments vary widely between each course. Registration in the majority of courses offered occurs only a few times over the entire period, see in Figure 1. The modules provided are updated each year, reflecting the changes to the course catalog via dropping courses and course code changes. Several introductory courses, required courses, and projects make up a significant portion of all enrollments. Most students at UCM need eighteen periods to complete their education. Nevertheless, some students enroll in over twenty periods. See Figure 2. Each recommender system focuses on a subset of twelve periods representing two years at UCM. The subset is refined further by selecting only students starting in the fall intake semester. These restrictions increase the standardization of students for our systems, and balance for the diversity of enrollment patterns present in an open-course curriculum. Our recommender systems use the remaining 1018 students that fall within these boundaries.

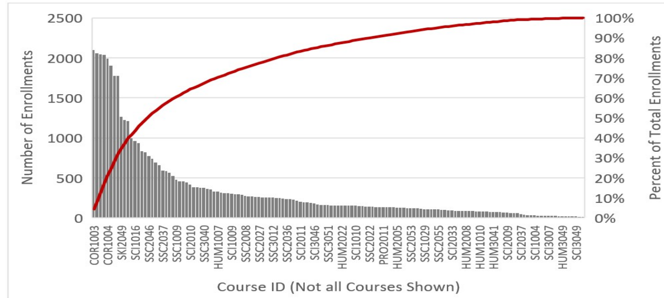

Figure 1: Total Number of Course Enrollments

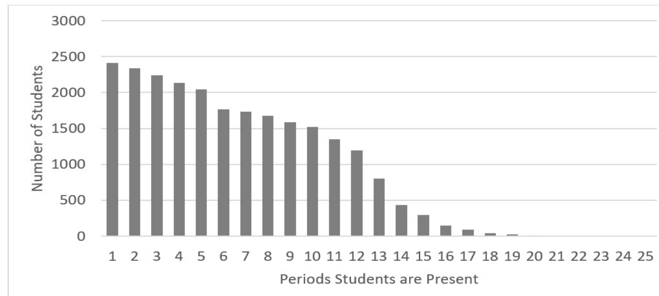

Figure 2: Total Number of Student Enrollments per Period

University College Maastricht offers a project-based curriculum. Two of the six periods each year are for student projects (periods three and six). The choice of courses within project periods is restricted. See Figure 3. Excluding these project periods, the average courses offered each period is 31 with a maximum of 44 courses. Figure 3 shows the variation in course offerings throughout the UCM data available. This variation is taken into account by our systems, and we omit these project periods from calibration. Course recommendations include only those courses offered in the target period. Therefore, recommending the maximum number of courses for a period results in a 0% error.

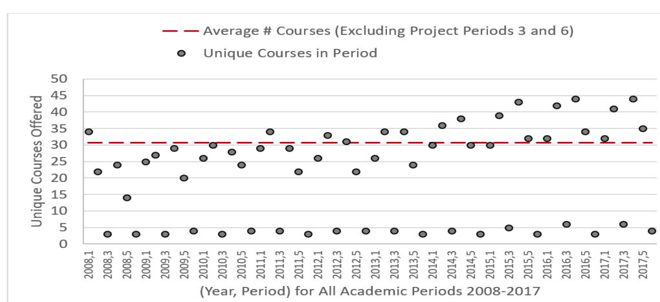

Figure 3: Courses Offered in Each Academic Period

7. EXPERIMENTS

This section presents experiments of ICCRS and ECCRS on the student-course data provided by University College Maastricht (see Section 6). First, an experimental setup is given, followed by results and discussion.

7.1 Setup

We validate ICCRS and ECCRS on the student-course data in the order of study periods. Assume that we have M number of periods $P_1, \ldots, P_M$ in which a student studies towards her degree (for our data M is equal to 18). We denote by $C_s(P_m) \subseteq C_s$ the set of courses that student s has taken in period P_m for $m < M$. Given new period P_{m+1} together with the set $\cup_{m=1}^{M} C_s(P_m)$ of courses student s has taken before that period, we test our recommender systems by checking whether the recommended sets C_s^ϵ of courses for P_{m+1} includes the courses $C_s(P_{m+1})$ that student s indeed has taken in P_{m+1}.

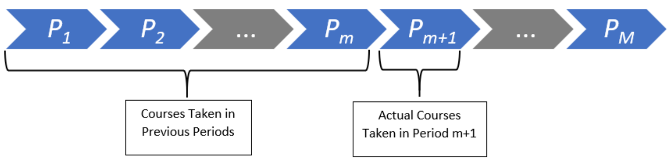

Figure 4: Prediction Process for Period P_{m+1}

We validate ICCRS and ECCRS using data from students in their second year of study. This choice is due to the fact the number of p-values possible is related to the number of courses a student has taken (see line 6 of the generic conformal course recommender in Algorithm 1). For example, a student with only three courses taken in period P_1 of year one can only have p values from the set $\{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ for any new course in period P_2.

For the validation process, we estimate the average error of e, and the average size of SR of the recommended sets C_s^ϵ of courses. We then use these statistics to study ICCRS and ECCRS as conformal predictors and as recommender systems.

In the first study, when we investigate ICCRS and ECCRS as conformal predictors, we are interested in establishing the validity and informational efficiency of the systems (check Sub-section 5.1). In the second study, when we investigate ICCRS and ECCRS as recommender systems, we are interested in estimating the error of the systems over the periods when we employ the recommended sets C_s^ϵ on a given course significance level ϵ. In our experiments, we use course significance levels ϵ of 0.05 and 0.1.

7.2 Results and Discussion

Figures 6 and 7 present the error plots and size plots of the recommended sets C_s^ϵ of ICCRS and ECCRS, respectively, for course significance level ϵ_c [2]. The error curves are very close to the diagonal $(0,0) - (1,1)$, which means the error is close to the course significance level ϵ_c. For ICCRS, the error is bounded mainly from above. For ECCRS, the error is bounded mainly from below. This bounding indicates that the systems are valid given sufficient information, especially ICCRS, which is conservatively valid [15].

The conservative validity of ICCRS explains why the averaged size SR of the recommended sets $C_s^{\epsilon_c}$ is higher than that of ECCRS. Thus, we may conclude that the informational efficiency of ECCRS is better in our experiments.

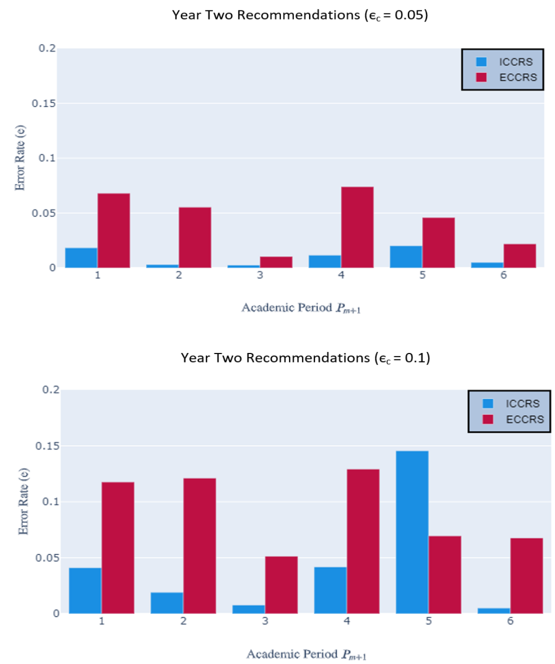

Figure 5: Period errors of ICCRS and ECCRS on course significance level ϵ_c of 0.05 and 0.1

Figure 5 presents the error of ICCRS and ECCRS when recommended sets C_s^ϵ on course significance levels ϵ_c of 0.05 and 0.1 are used. The systems are applied over periods of $P_1, P_2, P_3, P4, P_5$, and P_6 of the second year of the UCM students. The results show that:

- ICCRS and ECCRS produce accurate recommended sets of courses with an acceptable error.
- ICCRS is more accurate than ECCRS. This difference can be explained by the fact that ICCRS is more conservatively valid.
- the course significance level ϵ_c plays a substantial role: for 0.05, the error of recommended sets C_s^ϵ is much lower. However, this comes with a price: the size of the recommended sets is bigger when epsilon is lower.

8. CONCLUSION

This paper shows that safe course selection can be obtained if recommendations are supported with statistical confidence. The statistical confidence can be used for computing a statistically valid set of recommended courses that contains courses a student is likely to take with a probability of at least $1 - \epsilon$ for a user-defined significance level ϵ.

[2]We use the course significance level ϵ_c instead of the significance level ϵ for the predicted course sets since the range of ϵ is very restricted according to formula (3); e.g. $[0, 0.025]$ for 40 possible courses in a study period.

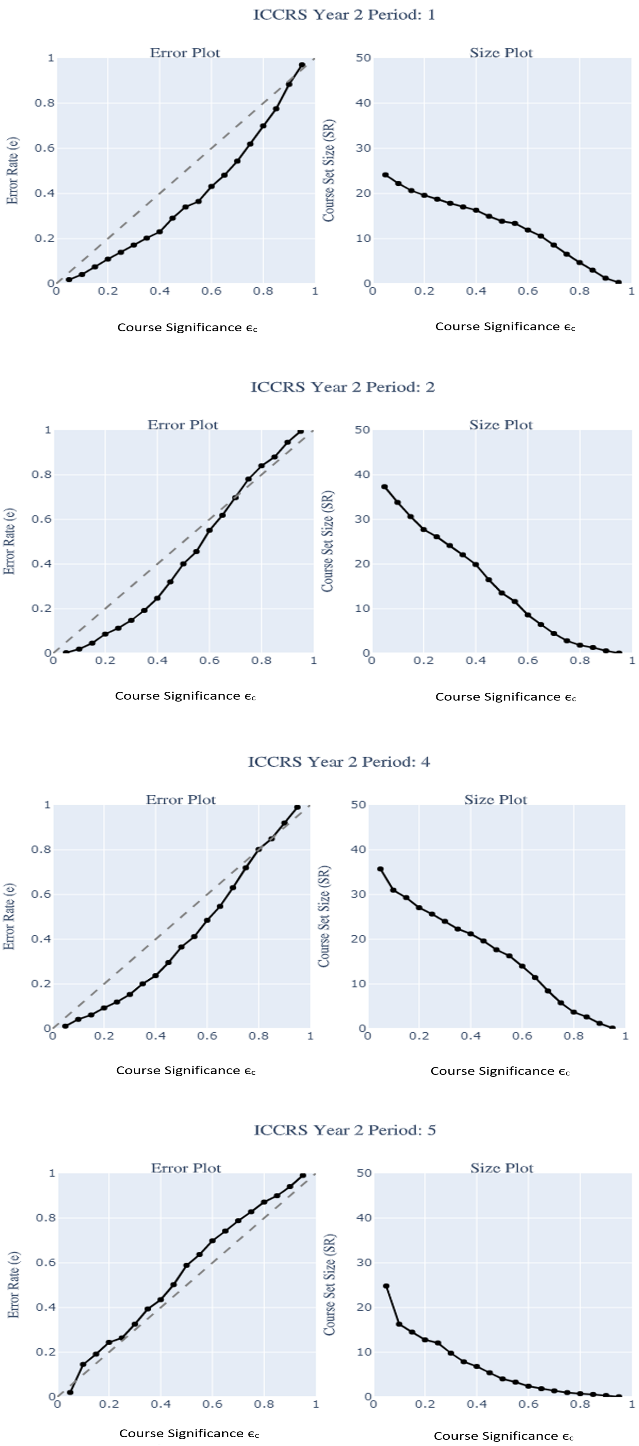

Figure 6: ICCRS - Calibration and Course Set Sizes

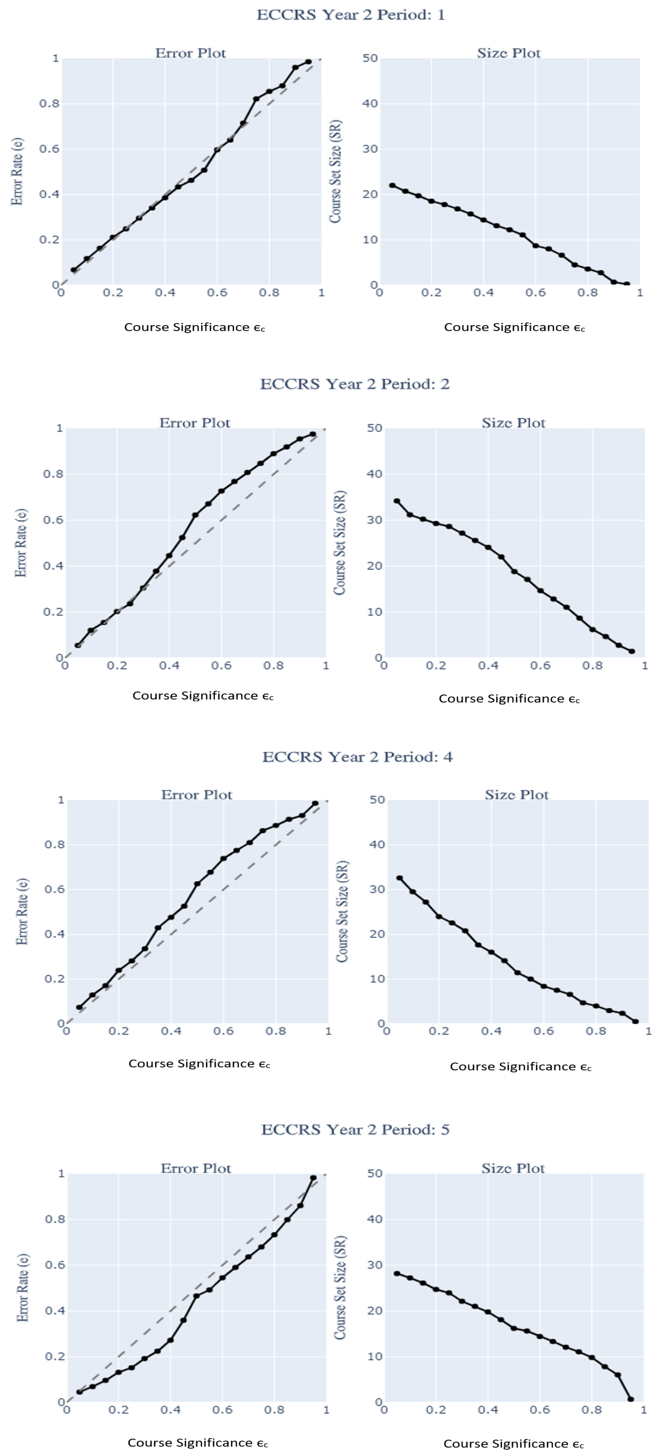

Figure 7: ECCRS - Calibration and Course Set Sizes

We have developed a generic conformal course recommender algorithm that outputs recommendations supported by statistical confidence. The algorithm has been instantiated in the form of two confidence-based course recommendation systems. The systems are essentially content-based: the first is an instance-based recommender system with relatively high accuracy. The second system is an exemplar-based system with a lower accuracy but with better explanatory capabilities. The experiments showed that both systems accurately suggest courses to students while providing statistically valid sets of courses recommended.

9. REFERENCES

[1] C. C. Aggarwal. *Recommender Systems: The Textbook*. Springer, 2016.

[2] D. Blei, A. Ng, and M. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

[3] H. Bydzovská. Course enrollment recommender system. In *Proceedings of the 9th International Conference on Educational Data Mining, EDM 2016*, pages 312–317. International Educational Data Mining Society (IEDMS), 2016.

[4] H. Dai, J. Liu, and E. Smirnov (eds.). *Reliable Knowledge Discovery*. Springer, 2012.

[5] K. Driessens, I. Koprinska, O. C. Santos, E. N. Smirnov, K. Yacef, and O. R. Zaïane. UMAP 2017 Education Recommender Systems Workshop. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization, UMAP 2017*. ACM, 2017.

[6] A. Elbadrawy and G. Karypis. Domain-aware grade prediction and top-n course recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*, page 183–190, 2016.

[7] T. Himabindu, V. Padmanabhan, and A. Pujari. Conformal matrix factorization based recommender system. *Information Sciences*, 467:695–707, 2018.

[8] V. Kagita, A. K. Pujari, V. Padmanabhan, S. Sahu, and V. Kumar. Conformal recommender system. *Information Sciences*, 405:157 – 174, 2017.

[9] A. Lambrou and H. Papadopoulos. Binary relevance multi-label conformal predictor. In *Proceedings of 5th International Symposium Conformal and Probabilistic Prediction with Applications, COPA 2016*, volume 9653 of *Lecture Notes in Computer Science*, pages 90–104. Springer, 2016.

[10] H. Ma, X. Wang, J. Hou, and Y. Lu. Course recommendation based on semantic similarity analysis. In *Proceedings of the 3rd IEEE International Conference on Control Science and Systems Engineering, ICCSSE 2017*, pages 638–641, 2017.

[11] R. Morsomme and S. V. Alferez. Content-based course recommender system for liberal arts education. In *Proceedings of the 12th International Conference on Educational Data Mining, EDM 2019*. International Educational Data Mining Society (IEDMS), 2019.

[12] M. P. O'Mahony and B. Smyth. A recommender system for on-line course enrolment: an initial study. In *Proceedings of the 2007 ACM Conference on Recommender Systems, RecSys 2007*, pages 133–136. ACM, 2007.

[13] M. Richter and R. Weber. *Case-Based Reasoning*. Springer, 2013.

[14] A. C. Rivera, M. Tapia-León, and S. Luján-Mora. Recommendation systems in education: A systematic mapping study. In *Proceedings of the International Conference on Information Technology Systems, ICITS 2018*, pages 937–947, 2018.

[15] G. Shafer and V. Vovk. A tutorial on conformal prediction. *J. Mach. Learn. Res.*, 9:371–421, 2008.

[16] E. N. Smirnov and A. Kaptein. Theoretical and experimental study of a meta-typicalness approach for reliable classification. In *Workshops Proceedings of the 6th IEEE International Conference on Data Mining, ICDM 2006*, pages 739–743. IEEE Computer Society, 2006.

[17] A. Surpatean, E. N. Smirnov, and N. Manie. Similarity functions for collaborative master recommendations. In *Proceedings of the 5th International Conference on Educational Data Mining, EDM 2012*, pages 230–231.

[18] A. Surpatean, E. N. Smirnov, and N. Manie. Master orientation tool. In *Proceedings of 20th European Conference on Artificial Intelligence, ECAI 2012*, pages 995–996. IOS Press, 2012.

[19] S. Zhou, E. N. Smirnov, and R. Peeters. Conformal region classification with instance-transfer boosting. *International Journal on Artificial Intelligence Tools*, 24(6):1–25, 2015.