

Can School Enrolment and Performance be Improved by Maximizing Students' Sense of Choice in Elective Subjects?

Rhyd Lewis^{1*}, Tom Anderson², and Fiona Carroll³.

Abstract

This paper explores a system that attempts to maximize high school students' sense of choice when selecting elective subjects. We propose that individual schools can tailor the combinations of subjects they offer in order to maximize the number of prospective students who can study their preferred subjects, potentially increasing enrolment numbers and academic outcomes while also reducing administrative overheads. We analyze the underlying computational problem encountered in this task and describe a suitable AI-based optimization algorithm that we have made available for free download. We also discuss some outcomes of using this method on a small number of case study schools.

Notes for Practice

- Schools can increase enrolment and income by strategically arranging the combinations of elective subjects on offer.
- This can be achieved using the free optimization software described in this article.

Keywords

Elective subject selection, student choice, optimization, graph colouring

Submitted: 05.04.2019 — **Accepted:** 10.01.2020 — **Published:** 03.04.2020

Corresponding author ¹ Email: LewisR9@cf.ac.uk Address: School of Mathematics, Cardiff University, Cardiff, Wales.

²Address: Bishop of Llandaff CiW High School, Cardiff, Wales.

³Address: School of Technologies, Cardiff Metropolitan University, Cardiff, Wales.

1. Introduction

Education today is not only about didactic instruction; it is about giving students the autonomy to take charge of their learning and empowering them to make choices about their educational pathways. Previous research has shown that “practices that facilitate and emphasise student choice — both in terms of the student’s educational path and the curriculum that the student follows — are widely considered to have positive impacts on student motivation and participation and, as a result, academic performance” (Hanover Research, 2014). As Little (1995) has commented, learners who accept responsibility for their learning are more likely to achieve their learning targets and, if these targets are met, will usually maintain a more positive attitude toward learning in the future.

This paper explores a system that maximizes students’ sense of choice regarding subject selection. Anders, Henderson, Moulton, and Sullivan (2017) have suggested that the subjects students choose to study at school will often have important consequences for their future academic and professional careers; however, in our experience it is not clear whether the correct combinations of subjects are always on offer. The preferred subjects of a student are often influenced by their personal characteristics and backgrounds, but they are also influenced by the school itself, with institutions often guiding students toward certain subject combinations simply because other combinations cannot be feasibly timetabled. Yamin-Ali (2014) has noted that limited subject choice is an issue facing students in schools in countries including Australia, Kenya, the United Kingdom, and Canada. In the same paper, a case study focusing on subject selection at a secondary school in Trinidad also found that (a) more than half of the form four students were not satisfied with their subject selection, (b) there was evidence of mismatch between students’ career preferences and their subject selection, and (c) students generally felt they needed more options in the subject offerings.

In many countries, students entering their final few years of study will be asked to choose a small set of elective subjects that

Table 1. Example Options Columns Used for Subject Selection (PE = Physical Education)

Please select one subject from four of the five columns below.				
Column 1	Column 2	Column 3	Column 4	Column 5
Maths	Further Maths	Art	Chemistry II	Biology
Media Studies	Geography	History	Computing	French
Chemistry I	Physics	PE	Spanish	

they wish to specialize in. In places such as the United States (APs), Germany (Abiturs), England and Wales (A-levels), and Australia (HSEs), these choices will be given in a set of “options columns” (sometimes also known as “blocking structures”), from which students are required to select a maximum of one subject per column. A small example is shown in Table 1. Note that when such structures are used, the choices that are available to students depend very much on the way that the subjects are partitioned. In Table 1, for example, students cannot study Further Maths with Physics because they are in the same column. On the other hand, they are free to study both Maths and Art, despite this being an unpopular subject combination in practice. Chemistry is also offered in two columns in this example because the school expects the number of students who choose it to exceed the maximum class size; hence multiple classes that cover the same material are necessary.

Schools use options columns in this way so that subjects from the same columns can be taught at the same time. So, using Table 1, we might choose to schedule all subjects from Column 1 during the first teaching period of the week, subjects from Column 2 in the second teaching period, and so on. Once these lessons (together with allocations of teaching staff) have been scheduled, schools can then set about timetabling the lessons of remaining year groups.

The way that subjects are split into options columns is important for schools and students for a number of different reasons.

- Our experience in the education sector indicates that in localities where schools compete for numbers, students are more likely to choose institutions that allow them to study their favoured options. Indeed, schools will often publicize their options columns online and during open days in order to entice prospective students. More students bring increased revenues. Although funding formulas can change, in the United Kingdom a single student who goes on to successfully complete their final years of study will typically earn between £2,000 and £5,000 of additional government investment for a state-funded school.
- Allowing students to select their favoured options is known to be beneficial for their studies. As Harackiewicz and Hulleman (2010) note while citing the work of Hidi (1990) and Schiefele, Krapp, and Winteler (1992), “being interested in a topic is a mental resource that enhances learning, which then leads to better performance and achievement.” From a school’s perspective, this is beneficial because it can bring higher rates of successful outcomes, which in turn boost outcome-related revenues. Anecdotally, we have also observed that allowing students to study their favoured subjects can improve student behaviour, attendance, and engagement, as well as reducing occurrences of students wanting to change their subject options mid-semester.
- In our experience, maximizing the number of students whose choices are satisfied by the options columns reduces the need for schools to be forced into creating additional classes (and incur the associated financial costs) in order to accommodate students whose choices are not met.

According to 2017 research carried out by two of England’s largest teaching unions, 64% of secondary school staff in England have observed a reduction in vocational subjects in their schools due to cost and performance pressures, leading to the limiting of student options (Burns, 2017). In Wales, the 2015 Donaldson Report advocated pupils having personalized “road maps” to avoid this reduction (Donaldson, 2015). Hypothetically, it is possible for schools to successfully pursue outcomes as their sole objective; small classes, personal tutoring for every pupil, and a large teaching staff would certainly deliver this. It is equally possible for schools to pursue austerity above all other matters; this would derive its success, conversely, by maximizing class sizes, with a skeleton teaching staff and every pupil forced to take the subjects that have the requisite gaps in their enrolment rosters. Clearly, there is a trade-off here, but opportunities for efficiencies should still be sought wherever possible.

In recent years in England, the number of students choosing art, languages, and social sciences at A-levels has fallen by around 10%. Meanwhile, entries in STEM subjects have increased by 15,500 students (6%) since 2013 (Reality Check team, BBC News, 2018). Clearly, student tastes and educational priorities are changing. Yet the combinations of subjects on offer at schools do not always reflect these changes, with many schools using the same, or only slightly modified, options columns year after year.

How can options columns be modified to fit the changing needs of students and schools? And how can individual schools tailor their options columns to meet the needs of their particular cohorts? Currently, many schools hold taster days and collect




Figure 1. Summary of the student subject preferences listed in Table 2.

data on the preferred subject combinations of prospective students. By knowing this information, schools can, in theory, ensure that their options columns allow the greatest number of students to pursue their favoured combination of subjects, enhancing the extent to which they feel enfranchised in their own education. However, we are unaware of any previous research or commercial software that has attempted to exploit such data. Perhaps the closest match is *SIMS Options* (2019), though this is more concerned with allowing students who are already enrolled at a school to select their subjects through an online system. Whether any optimization occurs behind the scenes, we are unable to ascertain.

In this paper, we propose that the task of producing options columns be modelled as an optimization problem. The aim in this problem is to maximize the number of prospective students whose preferred subject combinations are met entirely by a school’s options columns and, if this is not possible, to prioritize their favoured choices. In the next section, we begin by analyzing the computational problem underpinning this task and then discuss how to deal with additional constraints that are often imposed by schools. Section 3 describes our proposed AI-based optimization method. This software can be downloaded for free at <http://www.rhydlewislew.eu/resources/SubOpsCols.zip> and has been implemented in the VBA programming language, allowing it to be used within a simple Excel-based interface. Finally, in Section 4, we discuss some of our experiences of using this software in a small selection of schools, before making some suggestions for further work in this area.

2. Problem Formulation and Complexity

In this section, we analyze the problem of arranging subject options columns from a computational perspective. To aid these explanations, we will also use a worked example using a relatively small real-world data set collected at a taster day at a small private school in Cardiff, Wales.

To begin, the problem takes as input a list of subject preferences. This should be collected at the school by giving each prospective student a list of available subjects at the school and asking them to specify which ones they would choose to study if there were no restrictions in place. Information is also required regarding the maximum class sizes and the number of teachers available to teach each subject.

Our example input data is shown in Tables 2 and 3. It involves 45 prospective students and 20 subjects and requires a solution using four columns. We find that Biology and Chemistry is the most popular subject combination in this data (18 students), followed by Chemistry and Maths (15 students) and Biology and Maths (12 students). Biology is the most popular subject overall, with 27 students wanting to study it, though Maths and Chemistry are the most popular subjects as a first choice. Also note that Latin was not selected by any student and does not need to be considered further. Figures 1 and 2 summarize these features.

To find the best possible way of splitting subjects into columns, a naive algorithm might operate by simply going through every possible solution and figuring out which one satisfies the maximum number of student choices. In combinatorial mathematics, the number of ways of splitting n items into k subsets is given by a Stirling number of the second kind, denoted by $\left\{ \begin{matrix} n \\ k \end{matrix} \right\}$. Unfortunately, however, these numbers are known to grow exponentially with increases in n . For example, splitting our example data into four options columns gives $\left\{ \begin{matrix} 20 \\ 4 \end{matrix} \right\} = 45,232,115,901$ possibilities. For 30 subjects, this figure rises to $\left\{ \begin{matrix} 30 \\ 4 \end{matrix} \right\} = 48,004,081,105,038,305$ possibilities. Clearly, this sort of exhaustive approach is therefore only suitable for the very smallest of problem instances.

Table 2. Example Data: Subject Preferences for 45 Prospective Students Taken at a Taster Day in a Small Private School in Cardiff, Wales (RS = Religious Studies; DT = Design Technology; PE = Physical Education)

Student	Choice			
	1	2	3	4
1	Maths	Music	Physics	Eng Lit
2	Maths	Chemistry	Biology	Physics
3	Chemistry	Biology	Maths	PE
4	History	Economics	Eng Lit	Music
5	Maths	Economics	Gov & Politics	Psychology
6	Chemistry	Maths	Biology	Physics
7	Biology	Chemistry	Physics	History
8	Eng Lit	Gov & Politics	Psychology	Economics
9	Biology	Chemistry	Spanish	Physics
10	Eng Lit	Psychology	RS	History
11	Psychology	Spanish	Eng Lit	Biology
12	Economics	Maths	Gov & Politics	History
13	Eng Lit	Spanish	Psychology	Chemistry
14	Biology	Geography	Psychology	PE
15	Biology	Chemistry	Art	RS
16	Psychology	Biology	Eng Lit	-
17	Biology	Gov & Politics	RS	Psychology
18	Biology	Chemistry	Maths	Psychology
19	History	RS	Chemistry	Biology
20	History	Music	Maths	Spanish
21	Maths	Further Maths	Physics	Comp Sci
22	History	Eng Lit	Psychology	Gov & Politics
23	History	Biology	Gov & Politics	Spanish
24	Physics	Biology	Comp Sci	Psychology
25	Chemistry	Maths	Biology	-
26	Chemistry	Biology	Music	French
27	Maths	Further Maths	Physics	Chemistry
28	History	Physics	Maths	Geography
29	Biology	Physics	Comp Sci	-
30	Comp Sci	Maths	Physics	Chemistry
31	Chemistry	Biology	Maths	French
32	Chemistry	Biology	Maths	RS
33	Maths	Biology	Chemistry	RS
34	PE	Biology	Art	DT
35	Physics	History	Maths	Chemistry
36	Psychology	Economics	DT	-
37	Chemistry	Maths	Biology	Spanish
38	Psychology	Geography	Chemistry	-
39	Maths	Biology	Art	Chemistry
40	Maths	Chemistry	Biology	Physics
41	Chemistry	Biology	Physics	German
42	Maths	Further Maths	Physics	German
43	Art	RS	Eng Lit	Psychology
44	History	Economics	Psychology	Biology
45	Chemistry	Biology	Maths	Physics

Table 3. Maximum Class Sizes and Number of Available Teachers for Our Example Data

	Art	Biology	Chemistry	Comp Sci	DT	Economics	Eng Lit	French	Further Maths	Geography	German	Gov & Politics	History	Latin	Maths	Music	PE	Physics	Psychology	RS	Spanish
Maximum Class Size	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15
Available Teachers	1	2	1	1	1	1	1	1	1	1	1	1	1	1	2	1	1	2	1	1	1




Figure 2. Subject popularity and maximum class sizes for our example data.

2.1 Basic Formulation

A more suitable approach than exhaustive techniques is to employ some sort of combinatorial optimization algorithm. Given a list of students’ favoured choices, the first step in doing this is to arrange the input into an $(n \times n)$ matrix X , where the entry in the i th row and j th column (denoted by x_{ij}) gives the number of students who have stated a preference for studying both subjects i and j in the original data. In its simplest form, our task is to split the available subjects into k options columns so that the occurrence of popular subject combinations appearing in the same column is minimized. More formally, given a set of n subjects $\{1, \dots, n\}$, we want to partition this into k subsets $\mathcal{S} = (S_1, \dots, S_k)$ such that the following cost function is minimized:

$$f(\mathcal{S}) = \sum_{l=1}^k \sum_{i,j \in S_l} x_{ij}. \tag{1}$$

As shown, Equation (1) takes each column S_l of a solution in turn and considers all pairs of subjects i and j within it. If i and j are a pair of subjects that some students want to take, the value x_{ij} will be a positive value and will therefore contribute to the cost of the solution. Low-cost solutions therefore avoid situations where popular subject pairs appear in the same column.

Though this computational problem is quite simple to state, unfortunately it is not easy to solve to optimality. This is because it is equivalent to the well-known weighted graph-colouring problem, which belongs to a notorious family of problems known as the NP-complete problems.¹ Membership in this family implies that we cannot hope to find a fast-acting algorithm that will solve all instances of this problem to optimality. More formally, it tells us that there exists no exact algorithm whose time requirements are polynomially bounded according to problem size n (Garey & Johnson, 1979). This implies that we need to turn to more efficient but inexact heuristic methods to tackle this problem. Our suggested heuristic approach will be outlined in Section 3.

2.2 Partitioning Oversubscribed Subjects

Although Section 2.1 describes the core features of this problem, in practice the task of creating options columns for schools involves further considerations. The most pressing of these occurs when a subject is oversubscribed — that is, when the number of students who state a preference for studying a particular subject exceeds the maximum class size. In our example data, this occurs with four subjects, Physics, Maths, Chemistry, and Biology, as shown in Figure 2. As a result, two classes will be needed for each of these subjects.

Let us now consider one of these oversubscribed subjects, Chemistry. One option to use here (indeed, an option often used in schools) is to simply put Chemistry into one of the available options columns and then arbitrarily divide the students into two Chemistry classes at the start of the school year. However, it is likely that a greater number of student choices will be met by a solution if two instances of Chemistry are added to the solution in different columns, because students will then be able to choose the one that allows them to satisfy more of their other choices.

If multiple occurrences of a subject are required in a solution, each student will need to be assigned to one of these occurrences to allow us to construct the matrix X described in Section 2.1. Continuing with our Chemistry example, one approach would be to randomly partition the students into two classes, Chemistry I and Chemistry II, while ensuring the number of students per class is roughly equal. A better approach, however, is to partition the students so that those with similar choices in their remaining subjects are put in the same Chemistry class, because the remainder of the solution can then be arranged to accommodate these choices.

¹The equivalent “weighted graph-colouring problem” involves taking an edge-weighted graph G and partitioning the vertices of G into k “colour classes” such that Equation (1) is minimized, where x_{ij} gives the weight of the edge between vertices i and j . Further information on the graph-colouring problem can be found in the book by Lewis (2016).

To formulate this problem, imagine that s students have chosen to take a particular subject, and, as a result, c occurrences of this subject are required. Now let Y be an $(s \times s)$ matrix where element y_{ij} is the number of distinct subjects selected by students i and j . For example, from Table 2, we see that Student 2 has chosen Maths, Chemistry, Biology, and Physics, while Student 3 has chosen Chemistry, Biology, Maths, and PE. Therefore, the corresponding entry in Y will be five (Maths, Chemistry, Biology, Physics, and PE). Lower values in this matrix therefore indicate pairs of students whose subject choices are more similar to each other. Our task is to partition the s students into c subsets $\mathcal{T} = (T_1, \dots, T_c)$, minimizing the cost function

$$f(\mathcal{T}) = \sum_{l=1}^c \sum_{i,j \in T_l} y_{ij}. \quad (2)$$

Note that this function has the same characteristics as Equation (1); hence, this subproblem is also equivalent to the weighted graph-colouring problem, implying the need for heuristic algorithms. In our case, we require our algorithm to also be fast-acting (a) because the task of partitioning oversubscribed subjects can be considered a preprocessing step before the main optimization stage (outlined in Section 3) and (b) because it will need to run on each oversubscribed subject in turn.

As a result of these factors, our chosen method for this subproblem is based on the “steepest descent” search methodology. This operates by first producing an initial solution $\mathcal{T} = (T_1, \dots, T_c)$ randomly, while ensuring that the number of students in each subset T_i does not exceed this subject’s maximum capacity. In each iteration of the optimization process, this solution is then changed either by moving a student to a different group or by swapping the groups of two students. This change is chosen by looking at all possible moves and all possible swaps and carrying out the change that results in the biggest drop in cost according to Equation (2). The process repeats until none of these changes are seen to result in a cost reduction, meaning that the final solution corresponds to a local optimum in the cost landscape.

Table 4 shows a modified version of the data from Table 2 in which the required subjects — Chemistry, Physics, Biology and Maths — have each been partitioned in turn using this optimization process. If we again consider Chemistry, we see that 14 of the 23 students who chose this subject are now allocated to Chemistry I and that these students tend to have more similar subject choices. For example, of these 14, 12 have also selected Biology, and all of these have been allocated to Biology II. Similarly, all 14 of these students have also selected Maths, and all but 2 of these have been assigned to Maths II.

2.3 Further Considerations

In addition to the matter of partitioning oversubscribed subjects, consultations with our partner schools and others have also revealed a number of extra factors that need to be taken into account when constructing an options columns solution. In this section, we discuss how each of these can be incorporated into our problem model.

- Often, some pairs of subjects will be required to be in different options columns. This could be because the subjects need to be taught in the same room or by the same teacher, or because a student cannot select one subject without also taking another subject (such as Further Maths with Maths, or for a subject spanning multiple columns). If this is the case for two subjects i and j , then the elements x_{ij} and x_{ji} should be set to an arbitrarily high value in the matrix X . Any solutions featuring i and j in the same column will then incur very high costs according to Equation (1) and will be avoided by the optimization algorithm if possible.
- In a similar fashion, some subject pairs may need to be put in the same column, for example, to prevent students from taking both. In this case, the relevant elements in X should be set to some arbitrarily low (negative) value, encouraging the optimization algorithm to favour solutions in which these subjects appear in the same column.
- In some cases, schools prefer to produce options columns in which students’ first and second subject choices are prioritized over their remaining choices (some school administrators feel that prospective students will commit to a school when the options columns allow them to study their favourite two subjects only). This feature can be incorporated into the matrix by simply using larger values when counting first and second options. For example, in Table 4, History and Economics have been chosen as first and second choices on two occasions; however, instead of adding the value two to the relevant element in X , we might first multiply this by some positive factor. This same factor can then be used across all first and second choices.
- Finally, some schools or governments require that students study some mandatory subjects alongside their electives. For example, in New South Wales, students need to study English as part of their HSE qualifications; in other cases, religious schools might require students to take a theological subject as one of their electives. In our experience, schools cope with such requirements by first timetabling the necessary number of teaching sessions for the mandatory subjects. The remaining periods in the timetable are then used to accommodate the subjects listed in the options columns as described

Table 4. Our Example Input Once Students in Oversubscribed Subjects Have Been Partitioned

Student	Choice			
	1	2	3	4
1	Maths I	Music	Physics II	Eng Lit
2	Maths II	Chemistry I	Biology II	Physics I
3	Chemistry I	Biology II	Maths II	PE
4	History	Economics	Eng Lit	Music
5	Maths I	Economics	Gov & Politics	Psychology
6	Chemistry I	Maths II	Biology II	Physics I
7	Biology II	Chemistry II	Physics II	History
8	Eng Lit	Gov & Politics	Psychology	Economics
9	Biology II	Chemistry II	Spanish	Physics II
10	Eng Lit	Psychology	RS	History
11	Psychology	Spanish	Eng Lit	Biology I
12	Economics	Maths I	Gov & Politics	History
13	Eng Lit	Spanish	Psychology	Chemistry II
14	Biology I	Geography	Psychology	PE
15	Biology I	Chemistry II	Art	RS
16	Psychology	Biology I	Eng Lit	-
17	Biology I	Gov & Politics	RS	Psychology
18	Biology II	Chemistry I	Maths II	Psychology
19	History	RS	Chemistry II	Biology I
20	History	Music	Maths I	Spanish
21	Maths I	Further Maths	Physics I	Comp Sci
22	History	Eng Lit	Psychology	Gov & Politics
23	History	Biology I	Gov & Politics	Spanish
24	Physics II	Biology I	Comp Sci	Psychology
25	Chemistry I	Maths II	Biology II	-
26	Chemistry II	Biology I	Music	French
27	Maths I	Further Maths	Physics I	Chemistry I
28	History	Physics II	Maths I	Geography
29	Biology I	Physics II	Comp Sci	-
30	Comp Sci	Maths I	Physics I	Chemistry I
31	Chemistry I	Biology II	Maths II	French
32	Chemistry I	Biology II	Maths II	RS
33	Maths II	Biology II	Chemistry I	RS
34	PE	Biology I	Art	DT
35	Physics II	History	Maths I	Chemistry II
36	Psychology	Economics	DT	-
37	Chemistry I	Maths II	Biology II	Spanish
38	Psychology	Geography	Chemistry II	-
39	Maths II	Biology II	Art	Chemistry I
40	Maths II	Chemistry I	Biology II	Physics I
41	Chemistry II	Biology II	Physics II	German
42	Maths I	Further Maths	Physics I	German
43	Art	RS	Eng Lit	Psychology
44	History	Economics	Psychology	Biology I
45	Chemistry I	Biology II	Maths II	Physics I

in Section 1. Using our software, another option is to assign each prospective student to the mandatory subject(s) in the initial input data. These students will then be suitably partitioned into these subjects as described in Section 2.2.

Table 5 shows the final matrix X that was constructed from our example input. All choices are treated equally in this case. The dark shaded cells on the diagonal indicate the symmetry of the matrix. The lighter shaded cells containing an infinity value (∞) indicate pairs of subjects that should not be put in the same columns: Art and DT (because in this case the same staff member teaches both subjects) and the partitioned subjects (Biology, Chemistry, Maths, and Physics). All other cells in the matrix indicate the number of students who have been provisionally assigned to the corresponding pair of subjects. A zero entry here, such as between Art and Computer Science, indicates that no students have expressed a preference to study both subjects.

3. The Optimization Process

In this section, we describe how, having constructed our final matrix X , we can apply combinatorial optimization techniques to find high-quality solutions. In our case we use a technique known as simulated annealing. This is a general-purpose probabilistic

Table 5. The Matrix X Produced by the Example Data in Tables 2, 3, and 4; Values on a White Background Indicate the Number of Students from Table 4 Who Have Chosen Both Subjects

	1: Art	2: Biology I	3: Biology II	4: Chemistry I	5: Chemistry II	6: Comp Sci	7: DT	8: Economics	9: Eng Lit	10: French	11: Further Maths	12: Geography	13: German	14: Gov & Politics	15: History	16: Maths I	17: Maths II	18: Music	19: PE	20: Physics I	21: Physics II	22: Psychology	23: RS	24: Spanish
1: Art	0	2	1	1	1	0	8	0	1	0	0	0	0	0	0	1	0	1	0	0	1	2	0	
2: Biology I	2	0	8	0	3	2	1	1	2	1	0	1	0	2	3	0	0	1	2	0	2	6	3	2
3: Biology II	1	8	0	12	3	0	0	0	0	1	0	0	1	0	1	0	12	0	1	4	3	1	2	2
4: Chemistry I	1	0	12	0	8	1	0	0	0	1	1	0	0	0	0	2	12	0	1	6	0	1	2	1
5: Chemistry II	1	3	3	8	0	0	0	0	1	1	0	1	1	0	3	1	0	1	0	0	4	2	2	2
6: Comp Sci	0	2	0	1	0	0	0	0	0	0	1	0	0	0	0	2	0	0	0	2	2	1	0	0
7: DT	8	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1	0	0	1	0	0
8: Economics	0	1	0	0	0	0	1	0	2	0	0	0	0	3	3	2	0	1	0	0	0	4	0	0
9: Eng Lit	1	2	0	0	1	0	0	2	0	0	0	0	0	2	3	1	0	2	0	0	1	7	2	2
10: French	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0
11: Further Maths	0	0	0	1	0	1	0	0	0	0	0	0	1	0	0	3	0	0	0	3	0	0	0	0
12: Geography	0	1	0	0	1	0	0	0	0	0	0	0	0	0	1	1	0	0	1	0	1	2	0	0
13: German	0	0	1	0	1	0	0	0	0	0	1	0	0	0	0	1	0	0	0	1	1	0	0	0
14: Gov & Politics	0	2	0	0	0	0	0	3	2	0	0	0	0	0	3	2	0	0	0	0	0	4	1	1
15: History	0	3	1	0	3	0	0	3	3	0	0	1	0	3	0	4	0	2	0	0	3	3	2	2
16: Maths I	0	0	0	2	1	2	0	2	1	0	3	1	1	2	4	0	8	2	0	4	3	1	0	1
17: Maths II	1	0	12	12	0	0	0	0	0	1	0	0	0	0	0	8	0	0	1	4	0	1	2	1
18: Music	0	1	0	0	1	0	0	1	2	1	0	0	0	0	2	2	0	0	0	1	0	0	1	0
19: PE	1	2	1	1	0	0	1	0	0	0	0	1	0	0	0	1	0	0	0	0	0	1	0	0
20: Physics I	0	0	4	6	0	2	0	0	0	0	3	0	1	0	0	4	4	0	0	0	0	0	0	0
21: Physics II	0	2	3	0	4	2	0	0	1	0	0	1	1	0	3	3	0	1	0	8	0	1	0	1
22: Psychology	1	6	1	1	2	1	1	4	7	0	0	2	0	4	3	1	1	0	1	0	1	0	3	2
23: RS	2	3	2	2	2	0	0	0	2	0	0	0	0	1	2	0	2	0	0	0	0	3	0	0
24: Spanish	0	2	2	1	2	0	0	0	2	0	0	0	0	1	2	1	1	1	0	0	1	2	0	0

Table 6. Solution to Our Example Data Using $k = 4$ Options Columns

Column 1	Column 2	Column 3	Column 4
Biology I	Biology II	Art	Chemistry II
Chemistry I	DT	Comp Sci	Economics
Further Maths	Eng Lit	German	PE
Gov & Politics	French	History	Physics I
Music	Geography	Maths II	RS
Physics II	Maths I	Psychology	Spanish

Table 7. Solution to Our Example Data Using $k = 5$ Options Columns

Column 1	Column 2	Column 3	Column 4	Column 5
Economics	Art	Chemistry II	Biology I	Biology II
PE	French	Comp Sci	Chemistry I	DT
Physics I	Further Maths	Gov & Politics	German	Eng Lit
RS	Physics II	Maths II	History	Geography
Spanish	Psychology	Music		Maths I

SIMULATED-ANNEALING	
(1)	Produce an initial candidate solution \mathcal{S}
(2)	Choose an initial value for t
(3)	Set i to 0
(4)	while ($i \leq$ the iteration limit) do
(5)	Randomly modify \mathcal{S} to form a new solution \mathcal{S}'
(6)	if $f(\mathcal{S}') < f(\mathcal{S})$ then
(7)	Set \mathcal{S} to \mathcal{S}'
(8)	else with probability $\exp(-(f(\mathcal{S}) - f(\mathcal{S}'))/t)$
(9)	Set \mathcal{S} to \mathcal{S}'
(10)	Set $i = i + 1$
(11)	if i is a multiple of z then
(12)	Set $t = \alpha \times t$

Figure 3. The simulated annealing algorithm. Notation is described in the accompanying text.

optimization algorithm that is commonly used with this sort of logistical problem (Egeblad & Pisinger, 2009; Perea, Alcaca, Yepes, Gonzalez-Vidosa, & Hospitaler, 2008).

(As an aside, simulated annealing has also previously been used in a number of successful algorithms for timetabling lectures at universities (Cambazard, Hebrard, O’Sullivan, & Papadopoulos, 2012; Lewis & Thompson, 2015; Kostuch, 2005). This problem is related in part to the task of constructing options columns in that universities will often receive data in the form of a matrix X as above, with elements of X indicating the number of students who intend to enrol in the corresponding lectures. The process of dealing with this data at universities is different, however. Universities will typically seek to timetable lectures directly based on this matrix, with a number of different side constraints also being considered. These could include avoiding instances of students attending three or more lectures in consecutive time periods, ensuring students are granted lunch breaks, considering the availability of teaching staff, and making sure that lectures are scheduled in theatres with adequate facilities and seating capacities.)

The simulated annealing algorithm operates by taking a candidate solution and then, over time, making a series of small changes to this solution so that its cost decreases, thereby improving its quality. More specifically, at each iteration of the algorithm, the current solution, \mathcal{S} , is randomly modified to form a new solution, \mathcal{S}' . If this modification is seen to reduce the cost (i.e., $f(\mathcal{S}') < f(\mathcal{S})$), then \mathcal{S}' replaces \mathcal{S} as the current solution and the algorithm repeats. In addition, simulated annealing also allows modifications that worsen a solution to be performed from time to time to help the algorithm avoid getting stuck at local optima. Worsening modifications are accepted with a probability defined by the function $\exp(-(|f(\mathcal{S}) - f(\mathcal{S}')|)/t)$, where $|f(\mathcal{S}) - f(\mathcal{S}')|$ gives the proposed change in cost and t is a control parameter usually known as the *temperature*. Typically, t is set to a relatively high value at the beginning of execution, meaning that nearly all worsening modifications are accepted. During a run, t is then slowly reduced, meaning that accepting worsening modifications becomes increasingly less likely.

A pseudocode description of our simulated annealing approach is given in Figure 3. Here, the temperature t is reduced every z iterations by multiplying by a *cooling rate* α whose value should be set to between zero and one. Values of α closer to one cause the temperature to be reduced more slowly.

In our implementation, an initial solution \mathcal{S} is generated by randomly partitioning the subjects from the matrix X into k options columns, where k is defined by the user. At each iteration, a modification is then proposed either by moving a randomly selected subject to a new column or by swapping the columns of two randomly selected subjects. The cost of solutions is calculated using Equation (1). In addition, an initial value of 50 is used for t , together with an iteration limit of 1,500,000, a value for α of 0.995, and a z -value of 1000. Any elements in X containing the symbol ∞ or $-\infty$ are set to 1000 and -1000 , respectively, by default, though these can be changed by the user if desired.

3.1 Example Results

In this section, we look at some of the solutions produced by our methods using the example data from Table 5 as a small case study.

Table 6 shows a four-column solution. By cross-referencing with the data in Table 4, we find that 36 of the 45 students (80%) have all of their choices satisfied by this solution. Note that when an oversubscribed subject has been partitioned, a student does not necessarily need to select the one that they were provisionally assigned to in X . For example, in Table 4, we see that Student 1 has been assigned to Maths I, yet, in order to choose all of their four subjects, they must select Maths II instead. This demonstrates the advantages of assigning partitioned subjects to different options columns where possible. Of the nine students whose choices are not fully met by this solution, two cannot select their third options (Students 22 and 44) and seven cannot select their fourth options (Students 9, 10, 24, 26, 27, 28, and 43). Note also that, as required, Art and DT are assigned to different columns.




Figure 4. Profile of the simulated annealing execution that produced the solution given in Table 6. This run took approximately five seconds on a PC with a 3.3 GHz processor and 8 GB of RAM.

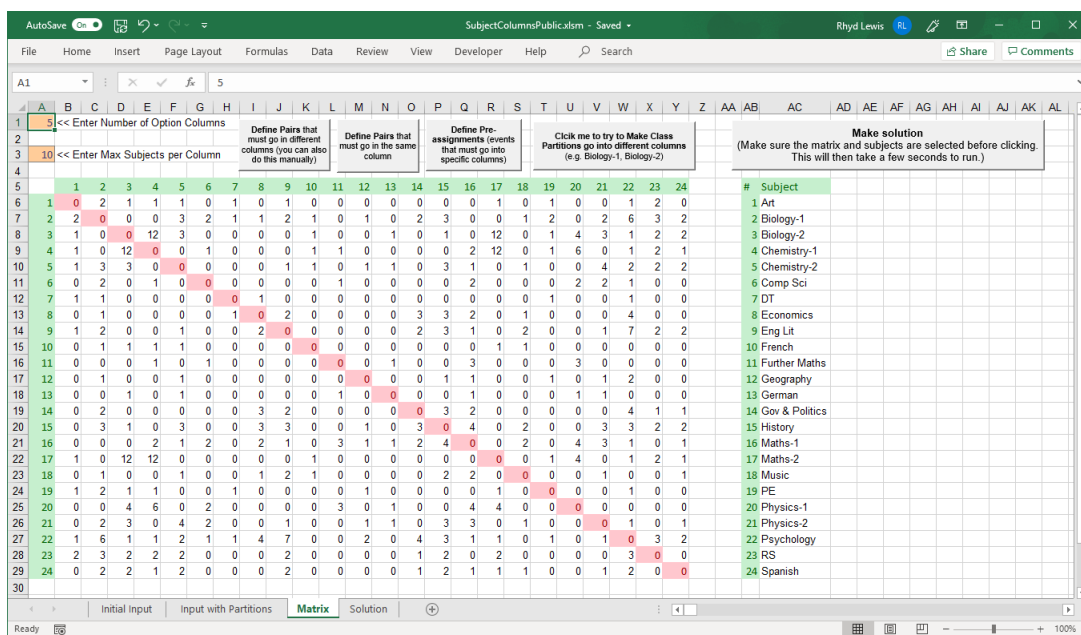


Figure 5. Screenshot showing an example matrix X formed by the input. Additional constraints can be added using the buttons provided.

Figure 4 shows how the simulated annealing algorithm can change the cost of the solution during execution. At the start, we see that the cost is quite high and fluctuating. This is because the initial solution is of poor quality and the temperature t is high, meaning that nearly all proposed modifications, improving and worsening, are accepted by the algorithm. However, as the run progresses and t is reduced, we see that the cost decreases and stabilizes. Such run profiles are typical of the simulated annealing technique (Kirkpatrick, Gelatt, & Vecchi, 1983). To contrast the quality of the solution returned in this run, we also ran our algorithm with this same data but deliberately left out the optimization steps used for partitioning subjects described in Section 2.2 (i.e., in this case students were randomly allocated to subject partitions). The resulting solution was able to meet the needs of only 28 of the 45 students.

One important feature of this problem is, of course, the number of options columns k permitted in a solution. At one extreme, if only one column is used, then solutions will only allow students to select one of their preferred subjects. On the other hand, if each subject is put into its own options column, then all students will be able to select all of their favoured subjects. Table 7 exemplifies this effect by showing a solution to our example data using five columns instead of four. In this solution, 43 of the 45 students (95.6%) now have their choices completely satisfied, with the remaining 2 (Students 1 and 43) being able to select everything except their fourth choice. Clearly, it is advisable for schools to increase their numbers of options columns where possible.

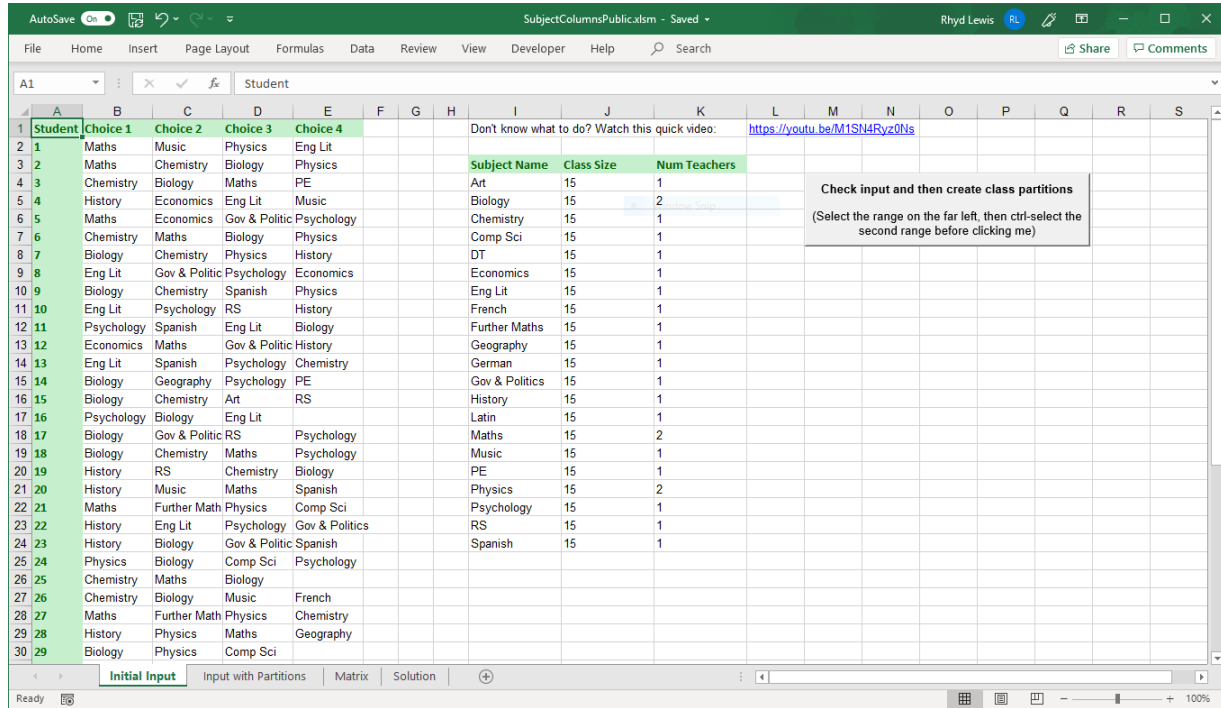


Figure 6. Screenshot showing how problem data is input into the software. Clicking on the button begins the interactive process.

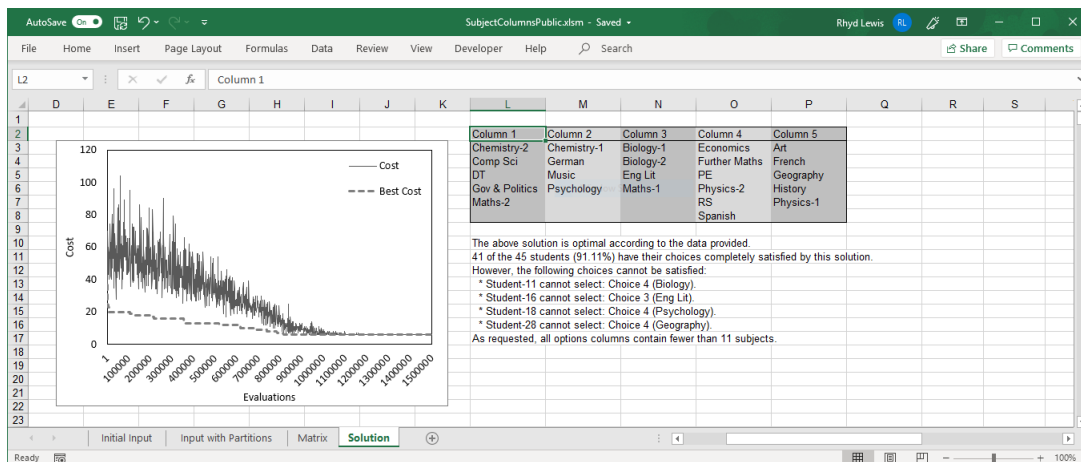


Figure 7. Screenshot of an example solution. Any problem requirements that are not satisfied by the solution are also listed by the software.

4. Outcomes and Discussion

For the past four years, we have been working with a number of partner schools in Cardiff, Wales, that encompass a variety of characteristics from the education sector. These include a high-performing state- and church-maintained school for 11- to 18-year-olds, the inner-city comprehensive school featured in the Channel 4 TV series *Educating Cardiff* (2019), and a smaller independent (private) school. According to consultations with the latter, we estimate that the administrative task of arranging a school's options columns can take an experienced member of staff up to 25 hours. In contrast, our software seems to bring the entire process down to around 30 minutes. Example screenshots of our software are shown in Figures 6 to 7.

At the school of this paper's second author, our optimization methods were first applied in 2016 using data collected from 16-year-olds about to embark on their A-level studies. That year, the school found that our solution resulted in a situation in which no incumbent student was forced to move to a different school due to a subject clash in the options columns. Meanwhile, more than 10 applicants from other schools cited the viability of the options columns as one of their main reasons for applying to the school. As a result, applications for the 2016 school year were seen to rise 32%, from 134 in 2015 to 177. In the following years, they rose further still to 193 in 2017 and 201 in 2018, bringing significant increases in revenue to the school. Alongside this was a steady upward trend in A-level results. When the 2016 cohort of students graduated in 2018, the number of A* and A grades rose 14% from the previous year to 38%, and, for the first time in the school's history, 100% of students achieved the Level 3 threshold² — an achievement that directly correlates with revenue levels under Welsh government funding formulas introduced in 2017–18.

At our inner-city partner school, prior to the use of our software, senior staff informed us that, for a number of years, they had observed a positive correlation between the quality of their options columns and pupil engagement (measured by how many pupils stayed at school, exclusions, attendance, behavioural incidents, and academic outcomes). In April 2017, we introduced our software, and the resulting solution was found to reduce the number of pupils who were unable to select more than two of their preferred subjects to 9, the lowest figure the school had ever had (previously 22). It will take more time to fully track the cohort who chose their subjects using this solution, but current projections anticipate an increase in all key indicators. This comes after a four-year downward trend, ending in August 2018, in which the Welsh education inspectorate ESTYN had graded the school as “needing improvement.” Senior managers at the school have confirmed to us that they believe these gains are partly due to the reduced disruption and improved teacher workload caused by fewer students changing subjects during the semester.

Although our empirical data is admittedly quite limited, we believe that these case studies serve to highlight the importance of options columns optimization in attracting students. As noted, we have also chosen to make our software freely available online so that other schools may use it and report their results to us. In the future, the ideas proposed in this paper might also bring business opportunities. Each school has a new cohort of students every year, and subject popularity inevitably fluctuates. As a result, schools have the opportunity to gain an advantage over their competitors by using options columns that maximally accommodate prospective students' needs. Another avenue to explore could be the licensing of the algorithm as a product sold either to consortia, to data-managing companies who assist schools (such as ALPS or SIMS), or simply to individual schools themselves. The use of this sort of software also allows the possibility of schools or consortia experimenting with various “what-if” scenarios. What would be the implications on student choice if one extra or one less column was permitted in a solution? What would happen if a school added a second instance of a subject to its teaching program? The effects of such changes on individual schools could easily be assessed by re-running the optimization algorithm with these different input sets.

Acknowledgements

The authors would like to thank Dave Roberts from Cathedral School, Cardiff, for useful discussions on producing options columns.

Declaration of Conflicting Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors declared no financial support for the research, authorship, and/or publication of this article.

²A government key performance indicator (KPI) that is satisfied when a student gains at least two A-levels with grades between A* and E.

References

- Anders, J., Henderson, M., Moulton, V., & Sullivan, A. (2017). *The role of schools in explaining individuals' subject choices at age 14* (Tech. Rep.). Centre for Longitudinal Studies, Institute of Education, University College London. <http://www.cls.ioe.ac.uk/anders>.
- Burns, J. (2017). *School budget squeeze "is reducing pupils' subject choice"*. <https://www.bbc.co.uk/news/education-39527183>. BBC.
- Cambazard, H., Hebrard, E., O'Sullivan, B., & Papadopoulos, A. (2012). Local search and constraint programming for the post enrolment-based timetabling problem. *Annals of Operational Research*, *194*, 111–135. doi: <https://dx.doi.org/10.1007/s10479-010-0737-7>
- Donaldson, G. (2015, February). *Successful futures: Independent review of curriculum and assessment arrangements in Wales*. <https://www.ewc.wales/site/index.php/en/documents/ite-accreditation/802-donaldson-report-successful-futures.html>.
- Educating Cardiff*. (2019). <https://www.channel4.com/programmes/educating-cardiff/>.
- Egeblad, J., & Pisinger, D. (2009). Heuristic approaches for the two- and three-dimensional knapsack packing problem. *Computers and Operational Research*, *36*(4), 1026–1049. doi: <https://doi.org/10.1016/j.cor.2007.12.004>
- Garey, M., & Johnson, D. (1979). *Computers and intractability — A guide to NP-completeness*. San Francisco: W. H. Freeman and Company.
- Hanover Research. (2014, November). *Impact of student choice and personalized learning*. <http://www.gssaweb.org/wp-content/uploads/2015/04/Impact-of-Student-Choice-and-Personalized-Learning-1.pdf>.
- Harackiewicz, J., & Hulleman, C. (2010). The importance of interest: The role of achievement goals and task values in promoting the development of interest. *Social and Personality Psychology Compass*, *4*(1), 42–52. doi: <https://doi.org/10.1111/j.1751-9004.2009.00207.x>
- Hidi, S. (1990). Interest and its contribution as a mental resource for learning. *Review of Educational Research*, *60*, 549–571. doi: <https://doi.org/10.3102%2F00346543060004549>
- Kirkpatrick, S., Gelatt, C., & Vecchi, M. (1983). Optimization by simulated annealing. *Science*, *4598*, 671–680. doi: <https://doi.org/10.1126/science.220.4598.671>
- Kostuch, P. (2005). The university course timetabling problem with a 3-phase approach. In E. Burke & M. Trick (Eds.), *Practice and theory of automated timetabling (PATAT) V* (Vol. 3616, pp. 109–125). Berlin: Springer-Verlag.
- Lewis, R. (2016). *A guide to graph colouring: Algorithms and applications*. Springer International Publishing. doi: <https://doi.org/10.1007/978-3-319-25730-3>
- Lewis, R., & Thompson, J. (2015). Analysing the effects of solution space connectivity with an effective metaheuristic for the course timetabling problem. *European Journal of Operational Research*, *240*, 637–648. doi: <https://doi.org/10.1016/j.ejor.2014.07.041>
- Little, D. (1995). Learning as dialogue: The dependence of learner autonomy on teacher autonomy. *System*, *23*(2), 175–181. doi: [https://doi.org/10.1016/0346-251X\(95\)00006-6](https://doi.org/10.1016/0346-251X(95)00006-6)
- Perea, C., Alcaca, J., Yepes, V., Gonzalez-Vidosa, F., & Hospitaler, A. (2008). Design of reinforced concrete bridge frames by heuristic optimization. *Advances in Engineering Software*, *39*(8), 676–688. doi: <https://doi.org/10.1016/j.advengsoft.2007.07.007>
- Reality Check team, BBC News. (2018). *A-levels: What subjects are students dropping and why?* <https://www.bbc.co.uk/news/uk-45171371>. BBC.
- Schiefele, U., Krapp, A., & Winteler, A. (1992). Interest as a predictor of academic achievement: A meta-analysis of research. In K. Renninger, S. Hidi, & A. Krapp (Eds.), *The role of interest in learning and development* (pp. 183–221). New York: Psychology Press.
- SIMS Options*. (2019). <https://www.capita-sims.co.uk/products-and-services/sims-options>.
- Yamin-Ali, J. (2014). *Data-driven decision-making in schools: Lessons from Trinidad*. Palgrave Macmillan.