Mathematical Techniques for Evaluating University Timetabling Policies

Matthew Davison[1,2], Ahmed Kheiri[3], and Konstantinos G. Zografos[4,5]

[1] Department of Mathematics, Aarhus University, Aarhus, Denmark
md@math.au.dk
[2] STOR-i Centre for Doctoral Training, Lancaster University, Lancaster, UK
[3] Alliance Manchester Business School, University of Manchester, Manchester, UK
ahmed.kheiri@manchester.ac.uk
[4] Edinburgh Business School, Heriot-Watt University, Edinburgh, UK
k.zografos@hw.ac.uk
[5] Distinguished Professor Emeritus, Department of Management Science, Lancaster University, Lancaster, UK

Abstract. University timetabling is a critical operational challenge, with policy decisions directly impacting staff, students, and resources. This paper presents a new benchmark instance derived from real-world data (*lancs-yr23*). Using this instance alongside an existing timetabling model and solution methods, we assess four policy questions: curriculum restructuring, enforcement of core teaching hours, relaxation of room capacity, and hybrid teaching under reduced physical capacity. This is done by modifying the model structure or data to represent each problem and solving the altered model. Our results show that students remain highly interconnected regardless of curriculum changes, that intuitive scheduling improvements can have counterproductive effects, and that modest capacity relaxation yields meaningful gains. *lancs-yr23* and the methodology provide a reusable framework for similar problems.

Keywords: Policy evaluation · Decision-making · Timetabling.

1 Introduction

The UK higher education (HE) sector generated £52 billion in total income in 2023/24 [22]. Tuition fees and education contracts account for over 52% of total income, dominating other income sources. Universities aim to attract as many students as possible to maximise this income.

A significant reason students choose a university is the quality of its teaching [2]. High-quality teaching, characterised by effective instruction and classroom management [6], requires knowledgeable staff and modern resources, both of which are costly. Despite the increasing income of the UK HE sector, Figure 1 shows that expenditure nearly matches and occasionally exceeds income. Data for 2023/24 shows that £18 billion went towards staff wages and £20 billion was spent on operating expenses [21]. This strongly motivates the need for careful staff and resource management.

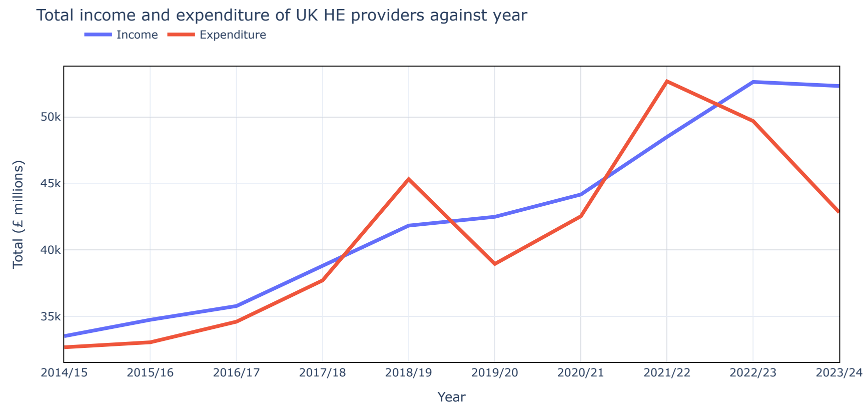

Fig. 1: The total income [22] and total expenditure [21] for higher education providers in the United Kingdom.

A central operational element at any university is the teaching timetable. This timetable dictates what staff, students and resources are doing at any given time. The construction of this timetable is limited by university policy. Consequently, modifying any university policy could have either a positive or detrimental impact on operational costs. This motivates the following research questions:

- Can existing timetables be used to identify specific policies problematic to university operation?
- Can changes in particular policies be modelled so that these changes can be incorporated into a timetabling model?
- Can the resulting timetables produced by this timetabling model be used to assess the overall impact of various policy changes?

To address these questions, this paper aims to: (i) reshape real-world timetabling data into a standard format, forming a novel benchmark instance, (ii) use this instance and knowledge from the university timetabling department to identify policies causing operational bottlenecks, and (iii) model and test changes to these policies, using change in timetable quality as a proxy for impact. Sections 2–5 cover related work, instance construction, policy evaluation, and conclusions, respectively.

2 Related work

For overviews of operational research in education, university course timetabling problem (UCTTP) methods, and benchmarking datasets, see [16], [5], and [4], respectively. Historically, timetabling was conducted manually from the prior schedule; [8] and [3] promote automatic methods to address the resulting inefficiencies. Building on this, researchers moved to developing models and methods that not only build schedules but act as tools for decision support.

The most direct approach is to build a new model tailored to the decision at hand. The model in [15] enabled the US Air Force Academy to switch from an alternating-day to a repeated-week pattern, yielding cost savings and class consolidations. [1] develop a mixed integer program (MIP) to assess a semester calendar change at MIT under COVID-induced resource scarcity.

Sensitivity analysis by modifying input parameters of these models can also produce insights. [23] use an integer programming timetabling model with varied input data to assess whether a schedule can be found if certain rooms are unavailable. Their case study helped identify teaching spaces the school could relinquish, albeit with some impact on timetable quality.

If a model already exists, some researchers modify or introduce constraints or objectives to represent new features. [7] optimise soft constraints relating to "compactness" (the arrangement of free time between activities), finding that their heuristic outperforms manual scheduling. [24] incorporate student flow modelling into an MIP for timetabling, significantly reducing corridor congestion. [18] modify the objective function of the timetabling problem described in [13] in various ways to incorporate "fairness". They then use this model to investigate the trade-off between timetable efficiency and fairness. [14] also pursue fairness in course timetabling via both modified objectives and additional constraints.

When several models share variables, these can be combined into a single monolithic model. [17] combine MIP models to explore trade-offs between different timetabling objectives.

This review highlights several ways practitioners can assess policy and other timetabling changes:

- Create a new model to replace outdated models.
- Sensitivity analysis by modifying input parameters.
- Modifying model aspects such as the objective function or constraints.
- Combine multiple models into a single model.

2.1 Contributions

The key gaps are the following: authors typically apply a single technique, give little guidance on which technique suits a given strategic problem, and seldom release case study data. This paper addresses all three, with the following contributions:

- A new benchmark instance for a large campus-based university in the United Kingdom and a thorough description of its creation.
- A demonstration of how to analyse the instance to identify bottlenecks caused by university policy or infrastructure.
- A study using the new instance in which we select and combine techniques to assess a range of policy changes, offering guidance on which technique suits which problem.

3 Instance construction and validation

The 2023/24 timetable at Lancaster University (LU) was provided by the LU timetabling team and exported from Scientia Syllabus Plus (SS+) as three files describing room, student and activity attributes. Some of this data was filtered in consultation with the timetabling team. Converting these exports into the format of the International Timetabling Competition 2019 (ITC-2019) [19] follows a largely routine extraction. We therefore restrict this section to the non-obvious decisions that depart from the ITC-2019 specification or that materially shaped the resulting benchmark.

3.1 Real-world data challenges

Two issues occur in real timetabling exports. Firstly, the timetable contains "tricks" used to force activities into the system which were removed on the advice of the timetabling team. For example, some field trips were booked as full-day campus events. Secondly, although no room may host two activities simultaneously in [11], the raw timetable violates this for a small number of spaces. Perfectly overlapping activities are treated as identical and merged, while partial overlaps are trimmed or split into parts, affecting under 0.02% of activities.

Room capacity was in the export, but availability was not, so rooms are assumed to be always available. Travel time within a building is taken as zero, and travel between buildings is computed from OpenStreetMap (OSM) data [20]. This is done by identifying the network node nearest to the centroid of each building, and the shortest-path time between nodes gives the building-to-building travel time (Figure 2). Spaces flagged with the "Lecture Capture or Stream" suitability are tagged as "hybrid capable", a property used in Section 4 and a room property not in the ITC-2019 specification.

Fig. 2: Map of the Lancaster campus at Lancaster University in the United Kingdom [20] with an example route between two buildings.

3.2 Adapting the ITC-2019 format to the model of [11]

The model of [11] shares many features with the ITC-2019 format, so most fields carry over directly; however, a few points need adaptation. The most substantial concerns module structure, which is inferred from the activities each student attends. Every distinct combination of activities forms a *module configuration*, with each activity a subpart. Because the raw data reuses activity identifiers across students, duplicates are relabelled to keep names unique and tied together by a "SameClass" constraint, equivalent to imposing the ITC-2019 "SameStart", "SameTime", "SameDays", "SameWeeks" and "SameRoom" constraints on the pair at once.

Time in the ITC-2019 is handled by describing each session with a start, a duration, and binary day and week patterns ("1001000" denoting Monday and Thursday); the set of times an activity may take forms its timeset. As a session carries a single start, duration and pattern, any activity whose sessions vary in these is split into separate sessions that a student must attend together. Rooms are assigned by giving each activity its largest suitable space or none (treated as online). Instructors shared across activities induce constraints that keep those activities from overlapping in time or sitting too far apart to travel between.

Finally, to give the optimiser room to manoeuvre in the experiments of Section 4, each activity is assigned candidate rooms (same zone, sufficient capacity, meeting all suitability requirements) and candidate timesets, generated by shifting its start, permuting its days, or permuting its active weeks while preserving weekly contact hours.

3.3 The *lancs-yr23* instance

The resulting instance, *lancs-yr23*, covers 65 weeks. On average, there are 7.39 modules per student and 5.74 activities per staff member. Other summary statistics are given in Table 1, and the instance is publicly available on the Lancaster University Research Directory [9].

Table 1: Summary metrics of the instance created using university data.

Metric	Activities	Modules	Spaces	Students	Timesets	Staff
Value	60,728	1,566	640	13,991	35,539	410

The LU 2023/24 timetable is a feasible solution to *lancs-yr23*, encoded in the instance as assignments with zero penalty. This is used as the initial solution throughout: both as the starting point required by the matheuristic of [10] and as the baseline against which policy changes are measured.

4 Policy evaluation

Using the techniques identified in Section 2, we assess changes to university timetabling policies by using *lancs-yr23* with some adaptation of the UCTTP model of [11] and solving or analysing the resulting model using the methodology of [10]. Each experiment modifies the model to capture the policy under study. We describe each modification in words and refer the reader to those works for the exact formulations.

Three objectives appear throughout this paper. z_1 is the number of elective module requests met (maximised); z_2 is the number of deviations from students' preferred teaching mode (minimised); and z_3 is the number of student scheduling issues (minimised). An elective module request is a module a student would like to take; teaching mode is how an activity is delivered, and a scheduling issue is a pair of a student's activities that overlap in time or are placed too far apart to allow travel between them.

The original timetable supplies a baseline, and unless stated otherwise, each experiment holds z_1, z_2 and z_3 to be no worse than their baseline values, isolating the effect of the policy change. After presenting each experiment, we review our findings and discuss the implications for strategic planning.

4.1 Curricula changes

Curriculum changes significantly affect timetabling, since university students often select which modules they take within their course. [10] proposes a measure for how "related" any two students are based on their module selection, together with a method for partitioning students on this measure. Using that methodology, we ask how related students are at present, how that relatedness changes under a modified curriculum, and the implications for the timetabling process.

In the raw LU data, the specific programme is given by the first four characters of the module, indicating the school or theme. For example, "MSCI222" is a module under "Management Science". The maximum number of schools/themes per student is eight. Let $n \in \{1, \ldots, 8\}$. For each n, an instance is created keeping only each student's top n schools/themes: if a student requests $\{\text{MSCI222}, \text{MSCI223}, \text{MATH240}\}$ and $n = 1$, the requests reduce to $\{\text{MSCI222}, \text{MSCI223}\}$, as "MSCI" occurs twice and "MATH" once. Students are then partitioned into module-independent blocks. Figure 3 shows how the number of blocks changes with n. The partitions created for $n \in \{3, \ldots, 8\}$ were all identical.

Table 2 shows that the largest partition block contains 91% and 93% of students for $n = 2$ and $n = 3$, respectively, indicating that any two students are likely connected by some chain of modules and that a clean partition is therefore difficult to achieve. Only when $n = 1$ does the method partition the students effectively, with the largest block in this case holding 6% of students.

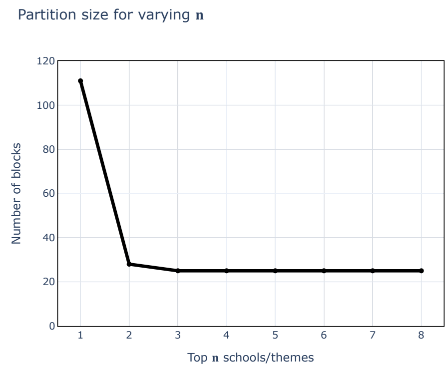

Fig. 3: Line plot illustrating the number of blocks in the partition when the top n schools/themes of each student are kept.

Table 2: Summary statistics for the partitions found.

Value of n	Block count	Largest block size	Smallest block size	Average block size
1	111	858	1	126
2	28	12,758	3	500
3	25	12,969	3	560

This experiment demonstrates that even radical curriculum restructuring leaves students highly interconnected through shared modules, so restricting interdisciplinary study would not meaningfully reduce timetabling complexity.

4.2 Teaching hours

At many universities, there are core teaching hours. At LU, these are 09:00 to 18:00 Monday to Friday, except Wednesday (09:00 to 13:00). The timetabling policy permits credit-bearing teaching outside these hours, but out-of-hours activities are disruptive: the Wednesday afternoon gap is reserved for extra-curricular activities, and early or late sessions disadvantage commuting students. In the baseline timetable, 3,814 activities (6% of the total) fall outside core hours, yet these affect 9,740 unique students (70%). Motivated by these figures, we ask whether activities can be moved within teaching hours and what the effect on infrastructure usage is.

One approach is to attach an integer penalty to each timeset equal to the number of meetings it places outside core hours. For a class held after 18:00 on a Monday and a Tuesday for three weeks, the penalty is

$$1\,(\text{event}) \times 2\,(\text{days}) \times 3\,(\text{weeks}) = 6. \tag{1}$$

We then minimise the sum of these penalties by applying the matheuristic of [10], which reduced the number of out-of-hours events from 3,814 to 3,751.

However, the new solution requires 35 more spaces (a 13% increase) and affects slightly more students (9,776 versus 9,740). This is a clear case of an intuitive improvement producing counterproductive results in practice.

4.3 Capacity relaxation

A key restriction at many universities is the availability of suitably sized teaching spaces. In the raw LU data, more than 400 activities had a planned or actual size exceeding the capacity of any suitable space. Figure 4 suggests why: many modules compete for the limited number of high-capacity venues on campus.

Room capacity is sometimes treated as a soft constraint in the literature [12], because attendance peaks at the start of a semester, drops, and then stabilises. The hard capacity constraints in this model create a tension between long-term efficiency and immediate student satisfaction. We therefore ask how many activities are over-subscribed, how necessary a one-to-one seating policy is, and whether a drop-off in attendance can be designed into the timetable.

Relaxing capacity first requires giving activities access to larger rooms, so we expand each activity's candidate locations to include rooms with enough capacity for its planned or actual attendance, whichever is smaller. Comparing each affected activity's attendance against its largest candidate room yields a set of *capacity deviations*. These distinct values, sorted into ascending order and prefixed with zero, form the relaxation levels $\mathcal{C}^{\text{dev}}$.

In [11], each activity's in-person attendance is bounded by the capacity of its assigned room. We relax this by granting each activity c a non-negative integer slack Δ_c above its room capacity, subject to a shared ceiling $\Delta_c \le \Delta$. Setting $\Delta = 0$ recovers the original hard constraint.

For each level $\delta \in \mathcal{C}^{\text{dev}}$, we fix $\Delta = \delta$, hold z_1, z_2 and z_3 to be no worse than baseline, and minimise z_2. Table 3 reports that relaxing capacity increases the number of mode preferences met. As Δ grows, more activities exceed their original capacity while the average relaxation used becomes a smaller fraction of the maximum permitted.

Table 3: Level of relaxation (Δ_c), total number of deviations from a student's preferred mode (z_2), number of activities with relaxed capacities, maximum relaxation and average relaxation.

	Δ_c			
	0	1	13	28
z_2	83,230	81,714	80,680	79,816
Relaxed capacities	0	303	501	655
Maximum relaxation	0	1	13	28
Average relaxation	0	1	5.54	8.04

Figure 5 shows how far each relaxed activity's attendance exceeds its room capacity. As the maximum permitted relaxation Δ increases, activities use a progressively smaller share of it, so most require only modest violations. A small Δ therefore secures much of the gain in mode preferences. Unlike a fully soft capacity constraint, this approach also bounds the worst-case overrun: a practitioner can read an acceptable Δ from such a histogram and continue optimising at that level.

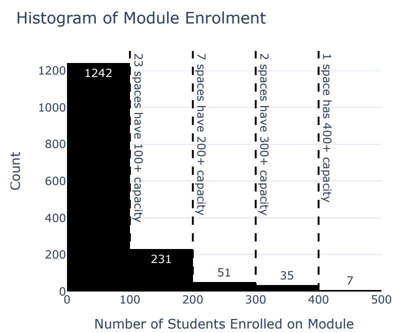

Fig. 4: Histogram showing the module enrolment figures. Information on physical space capacity is also given to highlight the issue of limited resources.

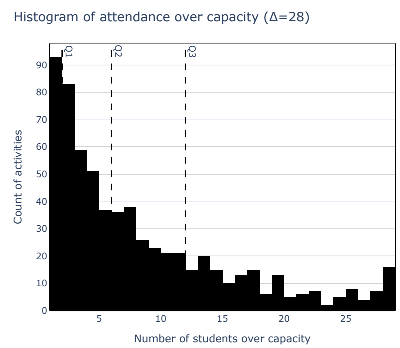

Fig. 5: Histogram showing the number of students over capacity for each relaxed activity. Quantiles are marked using dashed lines.

4.4 Hybrid teaching

Where the previous section addressed declining attendance, here we consider a drop in available physical capacity. During the 2020 COVID pandemic, university modules moved online to enable social distancing; on returning, some universities offered "hybrid teaching", in which part of a class attends in person and part online. Spreading students out for safe in-person attendance essentially reduced physical capacity. In some cases, this was a fourfold reduction [1]. Figure 6 shows that the baseline timetable becomes infeasible under such a reduction.

Some universities, including LU, have spaces that can record or broadcast lectures, though these are a small share of teaching spaces (19% at LU). We ask whether hybrid teaching resolves the capacity issue, whether the spaces to upgrade can be identified, and whether the number of upgrades can be minimised.

Solution under scarcity: Reducing every physical space's capacity by 75% makes the current solution infeasible, so everything is first moved online, and we then aim to return as many students as possible to in-person teaching. Concretely, we minimise the deviations from preferred mode (z_2) while holding the number of module requests met and the number of student conflicts to be no worse than their baseline values. Starting from an all-online solution avoids fixing the in-person/online split in advance, where a poor split could introduce conflicts through large travel times.

After some optimisation, only 50 activities returned to in-person teaching and 5 to hybrid mode. This limited gain could be a result of restrictive student timetables and the overall reduced capacity. As we are using a large-neighbourhood search, we are restricted to a limited number of iterations. Figure 7 shows that the objective does not plateau in the way a matheuristic near optimality usually does, suggesting that further iterations could restore more activities.

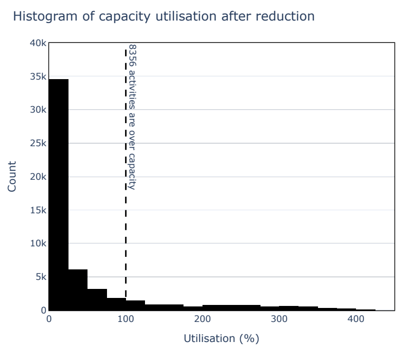

Fig. 6: Histogram showing theoretical utilisation of a space after reducing physical capacity by 75%.

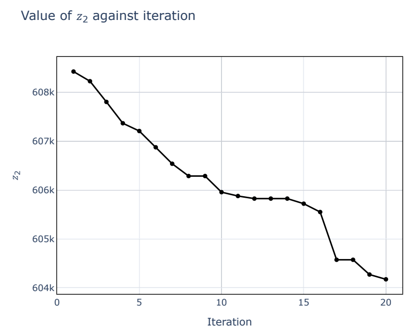

Fig. 7: Plot of the objective value against iteration in the case of restricted resources.

Solution allowing conversions: Not every space can host hybrid teaching, and conversions are costly, so we seek the fewest conversions necessary. We add a binary variable ∇_r marking each non-hybrid room r as converted (those in $R \backslash R^h$, where R^h is the set of hybrid-capable spaces), and relax the hybrid-compatibility constraint of [11] so a converted room may be used for hybrid teaching. We first allow conversions freely while improving z_2 and once this stalls, we fix z_2 and minimise the number of conversions,

$$z_5 = \sum_{r \in R \backslash R^h} \nabla_r. \tag{2}$$

Minimising z_5 identifies the essential conversions. Further optimisation improved z_2, returning 145 activities to in-person teaching (195 in total) and 26 to hybrid mode (31 in total). Post-optimisation analysis found that seven rooms had been converted for hybrid teaching, and we were unable to reduce this number by minimising z_5. Figure 8 shows the impact of each conversion on student and activity counts, identifying priority areas for hybrid enhancement. Spaces

113 and *253* each serve 11 activities and affect over 90 students, so they should be prioritised over the others.

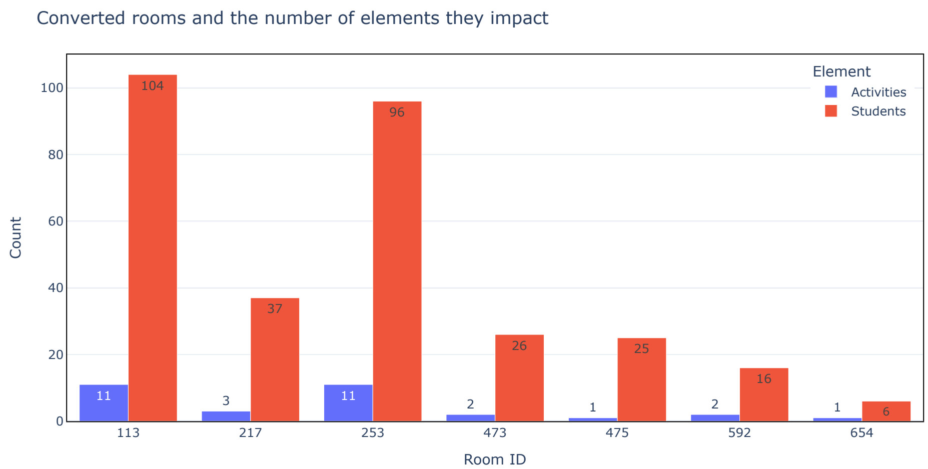

Fig. 8: Converted rooms and the number of activities using them, and the number of students this impacts.

5 Conclusion

This paper builds a new ITC-2019 benchmark instance from Lancaster University data, *lancs-yr23*, and uses it to assess four policy questions, drawing together the technique families of Section 2 in one case study. Three of the four build on the UCTTP model of [11]: the teaching-hours and capacity problems extend it with new objectives or constraints, and the hybrid-teaching problem combines several techniques at once. The curriculum problem instead modifies the input data and applies the partitioning scheme of [10], a form of sensitivity analysis. Three takeaways stand out:

- Building the benchmark was more than a format conversion: real timetables contain workarounds and overlapping bookings that must be cleaned up without distorting the problem.
- Intuitive fixes can make things worse. Moving teaching into core hours cut out-of-hours events but required more rooms and affected more students.
- The harder questions needed several techniques together rather than one. The hybrid-teaching study combined a model extension with a new objective to pinpoint the few rooms worth converting.

Future work could turn this into a catalogue pairing policy questions with the techniques best suited to them, helping practitioners choose an approach quickly. The *lancs-yr23* instance is publicly available to support this [9].

Acknowledgements We gratefully acknowledge the support of the EPSRC-funded EP/S022252/1 STOR-i Centre for Doctoral Training. We also thank the timetabling team at Lancaster University for providing us with timetable data and offering insight regarding timetabling challenges at LU.

References

1. Barnhart, C., Bertsimas, D., Delarue, A., Yan, J.: Course scheduling under sudden scarcity: Applications to pandemic planning. Manufacturing & Service Operations Management **24**(2), 727–745 (2022). https://doi.org/10.1287/msom.2021.0996, https://doi.org/10.1287/msom.2021.0996
2. Bhardwa, S.: Why do students go to university and how do they choose which one? Times Higher Education (Jun 2017), https://www.timeshighereducation.com/student/news/why-do-students-go-university-and-how-do-they-choose-which-one, accessed: 2025-06-06
3. Carter, M.W.: A comprehensive course timetabling and student scheduling system at the University of Waterloo. In: Burke, E., Erben, W. (eds.) Practice and Theory of Automated Timetabling III. pp. 64–82. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
4. Ceschia, S., Di Gaspero, L., Schaerf, A.: Educational timetabling: Problems, benchmarks, and state-of-the-art results. European Journal of Operational Research **308**(1), 1–18 (2023). https://doi.org/10.1016/j.ejor.2022.07.011, https://www.sciencedirect.com/science/article/pii/S0377221722005641
5. Chen, M.C., Sze, S.N., Goh, S.L., Sabar, N.R., Kendall, G.: A survey of university course timetabling problem: Perspectives, trends and opportunities. IEEE Access **9**, 106515–106529 (2021). https://doi.org/10.1109/ACCESS.2021.3100613
6. Coe, R., Aloisi, C., Higgins, S., Major, L.E.: What makes great teaching? review of the underpinning research. Tech. rep., Durham Research Methods Centre (DRMC) (2014), https://durham-repository.worktribe.com/output/1607293, full text deposited in DRO
7. Dammak, A., Elloumi, A., Kamoun, H., Ferland, J.: Course timetabling at a Tunisian university: A case study. Journal of Systems Science and Systems Engineering **17**, 334–352 (Sep 2008). https://doi.org/10.1007/s11518-008-5079-5
8. Daskalaki, S., Birbas, T., Housos, E.: An integer programming formulation for a case study in university timetabling. European Journal of Operational Research **153**(1), 117–135 (2004). https://doi.org/10.1016/S0377-2217(03)00103-6, https://www.sciencedirect.com/science/article/pii/S0377221703001036
9. Davison, M.: International Timetabling Competition 2019 approximation of Lancaster University 2023/2024 term. Dataset, file `lancs-yr23.xml` (Nov 2025). https://doi.org/10.17635/lancaster/researchdata/279, https://doi.org/10.17635/lancaster/researchdata/279, created 18 November 2025
10. Davison, M.: Modelling and Solving Multi-objective University Course Timetabling Problems. Ph.D. thesis, Lancaster University (2025). https://doi.org/10.17635/lancaster/thesis/3040
11. Davison, M., Kheiri, A., Zografos, K.G.: Modelling and solving the university course timetabling problem with hybrid teaching considerations. Journal of Scheduling **28**(2), 195–215 (Apr 2025). https://doi.org/10.1007/s10951-024-00817-w, https://doi.org/10.1007/s10951-024-00817-w

12. De Causmaecker, P., Demeester, P., Vanden Berghe, G.: A decomposed metaheuristic approach for a real-world university timetabling problem. European Journal of Operational Research **195**(1), 307–318 (May 2009). https://doi.org/10.1016/j.ejor.2008.01.043, https://www.sciencedirect.com/science/article/pii/S0377221708001665
13. Di Gaspero, L., McCollum, B., Schaerf, A.: The second International Timetabling Competition (ITC-2007): Curriculum-based course timetabling (Track 3). Tech. rep., University of Udine (2007)
14. Gogos, C., Dimitsas, A., Valouxis, C., Alefragis, P.: Modeling a balanced commute educational timetabling problem in the context of teaching integer programming. In: 2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM). pp. 1–5 (2022). https://doi.org/10.1109/SEEDA-CECNSM57760.2022.9932912
15. Gonzalez, G., Richards, C., Newman, A.: Optimal course scheduling for United States Air Force Academy cadets. Interfaces **48**(3), 217–234 (2018). https://doi.org/10.1287/inte.2017.0935
16. Johnes, J.: Operational Research in education. European Journal of Operational Research **243**(3), 683–696 (2015). https://doi.org/10.1016/j.ejor.2014.10.043, https://www.sciencedirect.com/science/article/pii/S0377221714008650
17. Lindahl, M., Mason, A.J., Stidsen, T., Sørensen, M.: A strategic view of university timetabling. European Journal of Operational Research **266**(1), 35–45 (2018). https://doi.org/10.1016/j.ejor.2017.09.022, https://www.sciencedirect.com/science/article/pii/S0377221717308433
18. Mühlenthaler, M., Wanka, R.: Fairness in academic course timetabling. Annals of Operations Research **239**(1), 171–188 (Apr 2016). https://doi.org/10.1007/s10479-014-1553-2, https://doi.org/10.1007/s10479-014-1553-2
19. Müller, T., Rudová, H., Müllerová, Z.: Real-world university course timetabling at the International Timetabling Competition 2019. Journal of Scheduling **28**(2), 247–267 (Apr 2025). https://doi.org/10.1007/s10951-023-00801-w, https://doi.org/10.1007/s10951-023-00801-w
20. OpenStreetMap contributors: Way: Lancaster University (23347159). https://www.openstreetmap.org/way/23347159 (2025), accessed: 2025-06-06
21. Perrott, L.: What is the expenditure of HE providers? Higher Education Statistics Agency (HESA) (May 2025), https://www.hesa.ac.uk/data-and-analysis/finances/expenditure
22. Perrott, L.: What is the income of HE providers? Higher Education Statistics Agency (HESA) (May 2025), https://www.hesa.ac.uk/data-and-analysis/finances/income
23. Schimmelpfeng, K., Helber, S.: Application of a real-world university-course timetabling model solved by integer programming. OR Spectrum **29**(4), 783–803 (Oct 2007). https://doi.org/10.1007/s00291-006-0074-z, https://doi.org/10.1007/s00291-006-0074-z
24. Vermuyten, H., Lemmens, S., Marques, I., Beliën, J.: Developing compact course timetables with optimized student flows. European Journal of Operational Research **251**(2), 651–661 (2016). https://doi.org/10.1016/j.ejor.2015.11.028, https://www.sciencedirect.com/science/article/pii/S0377221715010772