COURSE MANUAL
UNSUPERVISED & REINFORCEMENT MACHINE LEARNING

Research Master Business Data Science
Erasmus University Rotterdam, University of Amsterdam, Vrije Universiteit Amsterdam

Course Code	
Academic year	2019-2020
Period	1.3
Credits	4
Recommended knowledge	Business Foundations, Programming Basics (specifically programming in R, use of R Markdown or knitr), Mathematics, Statistics, Decision Theory for Business, Econometrics
Required knowledge	Linear Algebra, Linear and Logistic Regression

Table of Contents

1. COURSE COORDINATOR AND LECTURERS

Coordinator/Lecturer: prof. Gui Liberali (EUR)
Email: liberali@rsm.nl

Short bio: Gui Liberali is the Professor of Digital Marketing at the Erasmus University. His research has been published in Marketing Science, Management Science, IJRM, EJOR, and Sloan Management Review. He is the Vice-President for Membership at INFORMS Society of Marketing Science (ISMS), elected for the 2018/2019 term. He is currently co-editing a special issue of Management Science on Data-Driven Prescriptive Analytics. Gui is an ERIM Fellow(erim.nl) and twice was a finalist of the John Little award. His research focuses on multi-armed bandits, morphing theory and applications, and online experimentation. Gui was a visiting scholar at the MIT Sloan School of Management for several years. He holds a Doctorate in Marketing, and a B.Sc. in Computer Science

Lecturer: dr. Pieter Schoonees (EUR)
Email: schoonees@rsm.nl
Short bio: Pieter Schoonees is an assistant professor in the Department of Marketing Management at RSM, Erasmus University. His expertise lies in the fields of computational statistics, machine learning and psychometrics. A special interest is the use of such techniques for the analysis of data gathered from neuroscienfic studies.

2. COURSE CONTENT

Unsupervised & Reinforcement Machine Learning discusses unsupervised and reinforcement learning approaches often used to solve management science problems. The first part focusses on unsupervised machine learning techniques for finding meaningful relations between all variables in a data set simultaneously. In contrast to supervised machine learning, discussed in a previous course, in unsupervised techniques all variables play similar roles. Therefore, the relationships among all variables must be modelled, whereas in supervised learning only the relationships between the target variable and the features are of direct interest. An important application of unsupervised learning techniques in management is customer segmentation in targeted marketing.

The other main focus of this course is on reinforcement learning, where algorithms are used for sequential decision-making under uncertainty. Here, the focus will be on the design, solution, and implementation of reinforcement learning methods. Sequential decision problems involve a trade-off between exploitation (acting on the information already collected) and exploration (gathering more information). These problems arise in many important domains, ranging from online advertising, clinical trials, website optimization, marketing campaigning and revenue management.

An overview of techniques and ideas to be treated are:

- principal components analysis (PCA)
- cluster analysis (k-means, hierarchical)
- multidimensional scaling
- introduction to reinforcement learning and multi-armed bandits
    - Examples, formulation and preliminary results
- multi-armed bandit methods
    - Optimality of index-based policies
    - Heuristics: one-step look ahead, regret policies, Thompson sampling
- multi-armed bandit modeling strategies and applications.

3. LEARNING OBJECTIVES

By the end of the course students will be able to:

KNOWLEDGE AND UNDERSTANDING	Understand the fundamental building blocks of unsupervised and reinforcement machine learning methods, with special attention to: • Principal components analysis • Clustering techniques (k-means, hierarchical) • Multidimensional scaling • Index-based solutions to Multi-Armed Bandits • One-step look-ahead solutions to Multi-Armed Bandits

	• Thompson Sampling
APPLICATION OF KNOWLEDGE	Program these methods by translating technical knowledge of a method into their own code in R.
MAKING JUDGMENT	Apply and interpret these methods.
COMMUNICATION	Write a short report in the form of a scientific article.

4. STUDY MATERIAL

The course will cover material from the following list of readings, which are considered essential for your learning experience. These articles are also part of the examined material. Changes in the reading list will be communicated on Canvas.

Books:

- Hastie, T., Tibshirani, R. and J. Friedman. (2009). The elements of statistical learning (2nd edition). Springer. Available at https://web.stanford.edu/~hastie/Papers/ESLII.pdf.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). An introduction to statistical learning. Springer. Available at http://faculty.marshall.usc.edu/gareth-james/ISL/.
- Powell, W. and Ryzhov, I. (2012). "Optimal Learning". Wiley.
- Gittins, J.C.; Glazebrook, K, and R. Weber. (2011). Multi-armed bandit allocation indices. Wiley.

Selected papers, including:

- Groenen, P. J., & van de Velden, M. (2016). Multidimensional scaling by majorization: A review. *Journal of Statistical Software*, *73*(8), 1-26.
- Hauser, J.R., G. Liberali, and G.L. Urban. (2014). Website morphing 2.0: Switching costs, partial exposure, random exit, and when to morph. Management Science 60 (6): 1594–1616.
- Schwartz, E.M., Bradlow, E.T., and Fader, P.S. (2017). Consumer acquisition via display advertising using multi-armed bandit experiments. Marketing Science, 36(4).
- Scott, S.L. (2010). A modern Bayesian look at the multi-armed bandit. Applied Stochastic Models Business and Industry 26 (6): 639–658.
- Misra, K.; Schwartz, E., and Abernethy, J. (2019). Dynamic Online Pricing with Incomplete Information Using Multiarmed Bandit Experiments. Marketing Science, 38 (2): 226-252.
- Thompson, W (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika 25(3):285–294.
- Urban, G.; Liberali, G; Bordley, R.; Macdonald, E. and Hauser, J. (2014). Morphing Banner Advertising. Marketing Science, 33 (1), 27-46.

5. FORM OF TUITION

The lectures aim at stimulating your academic skills, and providing you with new knowledge. In this course, lectures are accompanied by tutorials. The information provided in the lecture is essential for the exercises, assignments and discussions during those sessions. We expect students to come to the lectures well prepared

and to participate in the interaction. Each week, one group will give a presentation on an exercise or assignment submitted before class.

The tutorials aim at practicing the theory using exercises and allowing students to ask for additional explanation for those parts of the material perceived as more difficult.

6. ASSESSMENT

Your overall course grade is composed based on different components that are presented in the assessment overview. You need a minimum overall grade of 5.50 to pass the course. The final grade is rounded to the nearest multiple of .0 or .5, with the following exceptions: any grade between 5.0 and 5.5 is rounded to a 5; a 5.5 is rounded to a 6; a 0.5 does not exist. Grades for homework or midterm examinations do not need to be rounded.

To prepare for the individual assignment at the end of the course, students practise in groups of three with similar group exercises and one group assignment. These assignments are aimed at being able to program a method in their own code and being able to write a report in article form. Feedback to exercises is given. The group assignment is graded. In each week, one of the groups also give a short presentation on the exercise or group assignment which is followed by a group discussion.

Format	% grade	Knowledge and Understanding	Application of knowledge	Communication	Learning Skills
Exam (Individual Assignment)	85 %	X	X	X	X
Group assignment	15%	X	X	X	X

Group Assignment - group assessment
During the first session, you will be assigned to a group of 3 students. The group assignments are designed to challenge you in various ways. They bring together different materials you study and practice during the lectures and tutorials, illustrating (some of) the principles learned in the course work and their applications (see **Appendix A**).

Individual Assignment
Your individual assignment is a paper in the format by the fast track of the Management Science journal. You will assessed based on the grades presented in **Appendix B**.

7. DETAILED COURSE SCHEDULE

Please check Canvas for an up-to-date schedule, reading material and assignments.

Week	Date	Time	Format	Theme/Topics	Preparation
1			Lecture 1 (2 times 2hr)	**Principal Component Analysis**	*Mandatory readings:* • Hastie et al. (2009), Chapter 14.5 • James et al. (2013), Chapter 10.2
			Tutorial 1 (1hr)	**Setting the Scene**	Managing expectations & Project Kick-Off • Team formation • Explain assignment 1 • Explain what we expect
			Submission Deadline	**Exercise Week 1**	Submit your exercise before Lecture 2 on Canvas (see Section 9.1)
2			Lecture 2 (2 times 2hr)	**Multidimensional scaling**	*Mandatory readings:* • Groenen & van de Velden (2016)
			Tutorial 2 (1hr)		
			Submission Deadline	**Exercise Week 2**	Submit your exercise before Lecture 3 on Canvas (see Section 9.2)
3			Lecture 3 (2 times 2hr)	**Cluster analysis (k-means, hierarchical)**	*Mandatory readings:* • Hastie et al. (2009), Chapter 14.3 • James et al. (2013), Chapter 10.3
			Tutorial 3 (1 hr)		Assistance for the group assignment
			Submission Deadline	**Group Assignment**	Submit your group assignment before Lecture 4 on Canvas (see Section 9.3)
4			Lecture 4 (2 times 2hr)	**Introduction to reinforcement learning and multi-armed bandits**	*Mandatory readings:* • Powell and Ryzhov, 2012 chapters 1,2 • Gittins and Glazebrook 1 and 9 • Urban et al. (2019)
			Tutorial 4 (1hr)		Assistance with Exercise Week 4
			Submission Deadline	**Exercise Week 4**	Submit your exercise before Lecture 4 on Canvas (see Section 9.4)
5			Lecture 5 (2 times 2hr)	**Dynamic Allocation Indices and One-Step Look Ahead Policies**	*Mandatory readings:* • Hauser et al., 2014 • Powell and Ryzhov, 2012 chapters 3, 5 and 6 • Gittins and Glazebrook 2, 3 and 6
			Tutorial 5 (1hr)		Assistance with Exercise Week 5
			Submission Deadline	**Exercise Week 5**	Submit your exercise before Lecture 5 on Canvas (see Section 9.5)
6			Lecture 6 (2 times 2hr)	**Thompson Sampling and other MAB heuristics**	*Mandatory readings:* • Misra et al., 2019 • Schwartz et al., 2017 • Scott, 2010 *Optional reading*: Thompson, 1933 Handing out the individual assignment (see section 9.6)
			Tutorial 6		Assistance with the individual assignment

Week	Date	Time	Format	Theme/Topics	Preparation
7			Exam		
	TBA		Submission deadline		Submit your individual assignment within two weeks after it has been handed out via Canvas (see Section 9.6)
	TBA		Exam Inspection	**Feedback**	Check Canvas for updates

8. ASSIGNMENT INFORMATION

9.1 Exercise Week 1

Write an article (maximum 4 pages) with introduction, data, methods, result, and conclusion sections. The central technique is principal components analysis. Provide your own code that computes a principal components solution and construct relevant graphics for interpretation. Compare the results of your code with that of the standard PCA implementation in R.

9.2 Exercise Week 2

Write the methods section for an article on multidimensional scaling (approximately 1 to 1.5 pages). Give an intuitive explanation in words what multidimensional scaling can be used for, provide technical details and diagnostics. Provide your own code for computing a multidimensional scaling solution using majorization and compare its results with a standard package in R.

9.3 Group Assignment Week 3

Write an article (maximum 4 pages) with introduction, data, methods, result, and conclusion sections. The central technique is k-means cluster analysis. Provide your own code that computes a cluster analysis solution and construct relevant graphics for interpretation. Compare the results of your code with that of the standard implementation in R.

9.4 Exercise Week 4

Write an article (maximum 4 pages) with introduction, data, methods, result, and conclusion sections. The central technique is Thompson Sampling. Provide your own code that a solution to a multi-armed bandit problem and construct relevant plots for interpretation.

9.5 Exercise Week 5

Write an article (maximum 4 pages) with introduction, data, methods, result, and conclusion sections. The central technique is a Dynamic Allocation Index such as Gittins Index. Provide your own code that a solution to a multi-armed bandit problem and construct relevant plots for interpretation.

9.6 Written Exam: Individual Assignment

The final exam is in the form of an individual take home assignment. It is handed out at the end of the final lecture in Week 6 of the course. The deadline is two weeks after handing out the individual assignment. The form of the individual assignment is the same as practiced in the group assignment and exercises: you will be assigned one or more methods and you will have to a report in the form of an article (consisting of introduction, data description, methods, results, and conclusion). In addition, you will have to program your own code for one of the methods in the course. Unlike the group assignments, here you have more pages to write (page limit is given in the Management Science fast track format).

APPENDIX A – INDIVIDUAL/GROUP ASSIGNMENT

Criterion	5 or lower	6	7	8	9 or 10
1. Research question **5%**	Question is unclear or illogical. Question is not functional. Question is too simple or too limited for the program or the study load.	Question refers to the technique, not to a substantive concept. Adequate and functional research question, but set at a minimum level of ambition.	Adequate and functional research question set at a level of ambition broadly appropriate for program and study load.	Well-formulated and clearly functional research question that can be answered by the available data and methods.	Original research question, displaying unusual insight and skill to translate relevant issues into well-formulated and researchable questions.
2. Method **40%**	Many of the elements under 7 are incorrect or missing.	As with 7 but some elements are incorrect or missing.	Methods section provides an intuitive explanation of the technique; the main goals and features are briefly described in simple terms. Notation is consistent and correct. The technical descriptions need to be correct. Key concepts of the method are correctly explained. Diagnostics of the method is properly described.	As with 7, explanation of techniques shows insight in the method.	As with 8, but addresses methodological and technical issues that go beyond what is covered in this program. Very extensive efforts in data / study collection.
3. Description and analysis of results **25%**	Many violations of the requirements outlined under 7. Poorly organized. Contains important errors of interpretation or logic; reveals lack of understanding of own research approach.	As with 7 but with several omissions. For example, standardized and/or mechanical presentation of results. Broadly effective, but inefficient or somewhat clumsy presentation. Contains minor errors of interpretation. Considerable unused potential for further analysis.	The order of discussion the results must be correct from more general modelling decisions towards interpretation. The results should be interpreted in terms of the meaning of the variables. The interpretation of the results should be sufficiently deep. Tables and figures support the decisions and interpretation, are readable, and have an informative caption.	As with 7 but is well-organized and thoughtful presentation of results, showing a good understanding of the nature of the data and many of the issues in interpretation.	As with 8, but very thorough analysis, showing a deep understanding of the research question, the research design, and the data. Presentation is highly effective.
4. Conclusion and discussion **5%**	No clear answer to research question, or an answer that does not follow from the research findings. No or just trivial suggestions for further research (e.g. 'collect more data').	Research question is answered by simple summary of findings. Perfunctory discussion of limitations and suggestions for further research.	Functional summary of findings, leading to discussion of extent to which research question is or is not answered. Meaningful reflection on limitations of own research. Some suggestions for useful further research.	A well-considered review of the findings in the light of the research question. Shows a clear understanding of limitations of own research. Several suggestions for further research that are properly explained.	Succeeds in putting the findings and the research question in the widest possible context, drawing out significant implications for theory development, research methodology and practice.

5. Written communication **10%**	Unstructured text. Fails to convey the key message of the report and/or to address questions. The text does not meet the academic editorial standards.	The text is somewhat unstructured and unclear. The text barely passes the academic editorial standards, as more polishing work is needed.	Overall well written, with occasional typos, or inaccuracies. The text passes the academic editorial standards, although the writing style is mechanical.	Structured text. The text is clear and concise, but here and there more (or less) details could improve the readability. Tables and Figures are self-explicatory and timely introduced in the text. The text meets the academic editorial standards, although the writing style is a bit mechanical at times.	Structured, coherent and polished text. Excellent writing style. The text is accurate, clear and concise, with the right level of detail. Tables and Figures are self-explicatory and timely introduced in the text. The text meets the academic editorial standards.
6. Code **15%**	Code is not working, difficult to read, or does not reproduce the results. Many of the requirements of 7 are missing.	As with 7, but some of the elements are missing.	Code implements the method correctly, is written for general data, structured, allows to reproduce the results. Naming of variables and comments are helpful to understand the code.	As with 7, but code is efficient and to the point.	Excellent code. Code goes contains an element of surprise that enhances its usage.
TOTAL					

APPENDIX B – EXAMPLE OF TAKE-HOME EXAM (INDIVIDUAL ASSIGNMENT)

Write a report of at most 4 pages (12pt font, no double columns) in which you solve a substantive research question using the techniques assigned to you. Use no more than 5 pages for introduction, data, methods, results, and conclusions, including tables and figures.

1. The report should be in the form of a small article (introduction, description of the data, substantive research question, methods, results, discussion and conclusions).
2. The description (including the technical properties, diagnostics, and usage) of the particular method is important.
3. Write a function for the technique assigned to you in R.
4. Use additional pages for code and the comparison of your program with the output of a standard function. Present your code as well as the script that runs the analysis. Make sure that your code can handle any data matrix of unknown size.
5. Search for your own data set from the sources indicated on the slides of week 6 or from another place. Describe the data briefly. Do not use a data set that have been used in this course before on the same technique by any of the groups or during the lecture.
6. Give a reference of the source of the data. If these data are your own, just say so.
7. Indicate whether the data propriety or can be freely used by others.
8. Do NOT copy and paste parts (e.g. method section) from your earlier assignments or exercises. Write them in your own words.
9. Substantive conclusions are important.
10. Show that you understand how to use the techniques sensibly.
11. Justify your conclusions by reporting appropriate results (possibly in tables or figures).

Deadline: two weeks after handing out the assignment. Check Canvas for definitive information.
Make a title page with: (in a large font) the number assigned to you, your name, your student number and your group number. (This page is not counted for the number of pages).