Learning Machine Learning: A Case Study

Niklas Lavesson

Abstract—**This correspondence reports on a case study conducted in the Master's-level Machine Learning (ML) course at Blekinge Institute of Technology, Sweden. The students participated in a self-assessment test and a diagnostic test of prerequisite subjects, and their results on these tests are correlated with their achievement of the course's learning objectives.**

Index Terms—**Artificial intelligence (AI), education, learning systems, prior knowledge, study results.**

I. INTRODUCTION

Students are attracted to courses in Machine Learning (ML) via their enthusiasm for Artificial Intelligence (AI), although the reality is that Machine Learning has become aligned with increasingly sophisticated mathematical and statistical techniques in recent years. At the same time, the analytical and mathematical skills of today's undergraduates are generally perceived as being low [1]. Several studies in higher education research report on the low student pass rate in mathematics. As a consequence, research into mathematics education at the undergraduate level has received great impetus in the last decade [2], for example through projects under the influence of the Bologna process in Europe. Recent years have shown an alarming drop in applications to computer science programs across Europe [3], and the students who apply do not seem to have as strong programming skills as before. Consequently, many universities seem to be lowering their standards to get a sufficient student quota. Some Swedish universities have decreased the level of mathematical experience needed to be admitted to engineering programs. The question is whether the fact that the students recruited seem to have less prior knowledge of computer science, mathematics, programming, and statistics affects the amount and depth of knowledge that can be learned. This correspondence addresses the question by investigating the relationship between prior knowledge and study results in the context of the topic of machine learning. The investigation was conducted by performing a case study on a group of students enrolled in an advanced-level Machine Learning course at Blekinge Institute of Technology (BTH) in Sweden.

A. The Role of Prior Knowledge

The increase in the number of people admitted to university studies has seemingly resulted in a greater variability in the background knowledge of first-year students. Consequently, it is argued that it has become even more important to understand the students' prior knowledge of a given subject [4]. If the relationship between prior knowledge, knowledge gain, and study results is understood, it might be possible to tailor the education to each student's needs and capabilities. An important question is how this tailoring would impact the learning process. Many studies address this question by investigating correlations between the prior knowledge and course results. However, intuitively, different subjects depend on prior knowledge to different extents. A certain subject might require generic knowledge (such as the ability to analyze), whereas another course might additionally require specific knowledge

(about, say, concepts and theories). The type of learning also varies extensively between different subjects and at different levels of education. Referring to Bloom's Taxonomy, learning can be classified according to the level of intellectual behavior involved. In contrast, it has also been shown that learning strategies and study patterns represent powerful factors related to the type of knowledge that is actually gained from studies [5]. The lack of some types of prior knowledge can be made up for by employing a productive learning strategy. That is, by adopting such an appropriate strategy, the student might be able to acquire the lacking prior knowledge in the process of studying the subject itself. A student's learning strategies and study patterns are based on his/her prior experiences of teaching and learning as well as on the design of the course and the written examination [6].

B. Related Work

Some studies have investigated the correlation between mathematical skills and student pass rates in informatics and computer science. For example, one study correlates student pass rates and student satisfaction (via course evaluation) with changes implemented in the Mathematics for Informatics course at the Faculty of Organization and Informatics in the University of Zagreb, Croatia [2]. A large number of changes were implemented in aspects such as content, learning and teaching methodologies, examination methods, and student support. The students were clearly more satisfied after the changes, and the pass rates were also considerably higher. Additionally, Lee *et al.* [4] present statistical regression models used for predicting students' first-year performance. According to this study, the results obtained from such models highlight that a mathematics diagnostic test is not only useful for gaining information on a student's prior knowledge, but is also one of the best predictors of his or her future performance. Several papers discuss the prerequisites needed to understand machine learning. For example, de la Higuera [7] argued that traditional theoretical computer science courses in graph theory, algorithm design, and complexity theory are needed. In contrast, Ambroladze and Shawe-Taylor [8] convey that modern machine learning is based on a broad spectrum of fundamental theories, from probability theory, through statistics and optimization, to information theory. There are quite a few related studies. However, most of these studies are position papers, and there seems to be only a small amount of research available on this critical issue.

II. BACKGROUND

Machine learning is the study of computer programs that improve automatically through experience [9]. It is generally regarded as a branch of computer science and more specifically as a subfield of AI. Much like AI, machine learning is a multidisciplinary research field, drawing from work conducted in fields such as statistics, mathematics, biology, control theory, philosophy, information theory, and psychology.

A. Machine Learning Research

The machine learning research field is closely related to pattern recognition and statistical inference [10]. It emerged in the late 1970s, and in 1980 the first machine learning workshop was held in Pittsburgh, PA. The first edition of the *Journal of Machine Learning* was published in 1986. From 1988 onward, the International Conference on Machine Learning has been held annually. Although the area of machine learning has been multidisciplinary in nature from its very inception [11], the current state of machine learning research is radically different from that of earlier research [12]. As an engineering field, machine learning has become steadily more mathematical and more successful in applications over the past 20 years [10]. In fact,

it is even argued that the increasing mathematical sophistication of machine learning in recent years has reached a level at which some premium conferences deter all but the most mathematically sophisticated researchers [1]. Conversely, the state-of-the-art learning methods (algorithms) are based on complex and advanced concepts, such as quadratic programming, probabilistic graphic models, hidden Markov models, and entropy measures. The learning techniques and the theoretical foundations behind them are becoming more and more complex, as are the applications in which they are used. Thus, in order to understand machine learning, a novice must at the very least have a basic understanding of background subjects, including linear algebra, basic mathematics and calculus, statistics and probability, and programming. Arguably, without this proper background, the subject can only be understood at a general and rather simplistic level.

B. Machine Learning Education

According to the Bologna process, higher education is divided into three levels: basic, advanced, and research studies. Machine learning courses are taught at various levels of education. However, they seem to be predominantly given at the advanced or research level of computer science programs [1]. The prerequisites for admission to Master's-level computer science programs vary between universities. At BTH in Sweden, the prerequisite is a B.Sc. degree in computer science or a corresponding level of knowledge. The requirements for receiving a B.Sc. degree are a total of 180 ECTS credit points, including 90 in computer science and at least 15 in mathematics. For the purposes of this correspondence, it is sufficient to know that 1.5 ECTS credit points roughly correspond to one week of full-time studies (40 h). Thus, 180 ECTS credit points represent three years of full-time studies. The majority of computer science courses at BTH are practically oriented. Thus, besides the theoretical foundations, most exercises and assignments include various forms of software design, documentation, and programming. However, the majority of the Master's students at BTH originate from outside Europe and tend to have a more theoretical background in computer science, although experience has shown that this theoretical background, as well as the non-European students' learning strategies and study patterns, seem to be vastly different to those of the domestic students.

The Machine Learning course at BTH is given at the advanced level and is featured in a number of computer science Master's programs. The course comprises 7.5 ECTS credit points and corresponds roughly to five weeks of full-time studies. While the course has a strong focus toward supervised learning and data mining, the curriculum also includes reinforcement learning and unsupervised learning.

According to the aims and learning outcomes that are stated in the course descriptor, at the end of the course the student should be able to:

LO1 Independently and thoroughly evaluate and compare the performance or, other qualities, of algorithms for typical learning problems;

LO2 In a group, or independently, implement learning algorithms on the basis of algorithm pseudocode and scientific papers or books;

LO3 Independently and thoroughly describe and compare different evaluation methods for learning algorithms;

LO4 Independently and briefly describe and compare different machine learning paradigms;

LO5 In a group, or independently, plan data mining experiments, apply data mining tools to run the experiment, and finally, gather information from the experimental run and present this in a report.

At BTH, students are expected to learn about and make use of an open-source machine learning workbench, called Weka [13], to conduct experiments and solve assignments. Basically, Weka is a piece of software that can be used to analyze data sets, perform machine learning and data mining experiments, and analyze experimental results. However, Weka may also be used as an application programming interface (API) to develop revised or new learning algorithms as well as complete applications. In order to use Weka as an API, the student has to have at least basic programming skills. The mandatory assignments of the BTH Machine Learning course include performing an experiment using Weka, developing a simple supervised learning algorithm, and developing a reinforcement learning application.

III. Method

The focus of the study presented here is a group of $n = 11$ respondents from a class of approximately 30 students enrolled in the Machine Learning course at BTH. The aim of this case study is to investigate whether there is a correlation between prior knowledge of statistics, mathematics, and programming, and the eventual examination results. The case study comprises four different stages of data collection. During the fifth out of a total of 10 lectures, the students are asked if they would like to participate in taking a test whose objective is to gather information that can be used to improve the learning of future students by modifying the presentation and contents of the course.

The students who do agree to take this test are given a list of five subjects for which they should rate their own level of knowledge. After completing the self-assessment test, the students are given 10 min to participate in a diagnostic test that features two elementary questions from each of the subjects. When the final written examination has been held and the course has officially ended, the results for the three assignments and the written examination are collected from the Student Documentation System (Ladok) in order to be used for analysis.

A. Self-Assessment

At the self-assessment stage, the students are asked to estimate their own level of knowledge of five subjects from 1 (no knowledge) to 9 (extensive knowledge). The featured subjects are probability and statistics, basic mathematics and calculus, linear algebra, discrete mathematics, and programming. The rationale for selecting these particular subjects is that their theory and concepts form the basis of state-of-the-art learning algorithms and are used in common machine learning textbooks [9], [13] to explain learning theory and to describe algorithms.

B. Diagnostic Test

The diagnostic test features two elementary questions for each subject. It is difficult to measure the knowledge of a subject with only two short questions. However, an extensive set of questions for each subject might make some students more reluctant to participate. In addition, it is hard to make an unbiased selection of questions, that is a selection that properly represents the basic knowledge of a subject. It was therefore decided that questions would be derived from the introductory examples found in the aforementioned textbooks. The students were given 10 min to answer the questions below, which are designed to elicit the responder's basic knowledge of each subject.

1) Probability and Statistics
    a If a fair dice is thrown twice, what is the probability that six turns up both times?
    b If the mean of a sample is 90 and the variance is 0.25, how big is the standard deviation?
2) Basic Mathematics and Calculus
    a Find the derivative of $f(x) = x^2 + 3x$.
    b If $a = 5$ and $b = 2$, what is the result of the calculation of $(a - b)b$?

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

IEEE TRANSACTIONS ON EDUCATION 3

3) Discrete Mathematics
   a If we have two sets, $X = \{2, 3, 4\}$ and $Y = \{2, 1, 5, 4\}$, how many elements does the union of X and Y have?
   b If we have a set, $X = \{1, 2, 3, 4\}$, what is the absolute, $|X|$, of this set?
4) Linear Algebra
   a If we have a line, $y = 5x + 2$, how much does y increase for each x-step and at which y does it intersect the x-axis?
   b Calculate the Euclidean distance between the points $(0, 0)$ and $(3, 4)$.
5) Programming
   a What is the output of the following code?
   b What is the output if % is changed to /?

```
for (integer i = 0; i < 4; i++)

{

if (i % 2 == 0) Print(i);

}
```

C. Course Results

The students' results are extracted from Ladok, which is the national system used for documentation of academic information at higher education institutions in Sweden. As stated earlier, the Machine Learning course at BTH comprises 7.5 credit points. It features two small assignments (1 credit point each), one larger assignment (2 credit points), and a written examination (3.5 credit points). The ECTS-based course grade is calculated from the accumulated score on the written examination. The self-assessment test is entirely subjective, whereas the diagnostic test and the course results provide objective measures of the students' knowledge of mathematics, programming, and machine learning. However, there are several issues to consider. First, the written examination should correspond to the learning objectives, as specified in the Bologna-based course descriptor. It is difficult to evaluate objectively whether the written examination questions actually do correspond to the learning outcomes or not. Second, the learning objectives should represent a basic understanding of the topic of machine learning. This is perhaps even more difficult to assess. To some extent, it is possible to use information obtained from the results of the course evaluation in order to understand, from a student's perspective, whether the expectations raised from the course descriptor are indeed met.

D. Course Evaluation

BTH uses a Web-based course evaluation questionnaire, which provides a means for the students to voice their opinion on the course and to rate components such as teaching capabilities, level of difficulty, and content. The results from the course evaluation are obtained shortly after the completion of the course. Parts of this information may be used as a basis for discussion and better understanding of the test results. The Web-based questionnaire consists of four sections: classification, course experience, learning objectives, and suggestions for improvement. At BTH, the course experience is evaluated using the Student Course Experience Questionnaire (SCEQ) [14], developed at the University of Sydney, Australia. The purpose of the SCEQ is to provide the university community with a basis for strategic faculty-level academic development and curriculum review to improve the quality of teaching and student learning. Among other things, the BTH system extends the SCEQ by providing a way for the students to assess subjectively their own fulfillment of the learning objectives. Additionally, it was decided to perform interviews with the participating students in order to get a deeper understanding of the problem domain.

TABLE I
SELF-ASSESSMENT AND DIAGNOSTIC TEST RESULTS

Subject	Assessment[1] mean (SD)	Correct Answers[2]
Probability and Statistics	5.91 (1.51)	.27
Basic Mathematics & Calculus	7.27 (1.01)	.73
Discrete Mathematics	6.27 (1.27)	.41
Linear Algebra	6.82 (1.47)	.23
Programming	6.73 (1.95)	.50

[1] The mean self-assessment scores for the five included subjects
[2] The overall ratio of correctly answered questions on the diagnostic test

IV. RESULTS

The Machine Learning course is held once every fall at BTH; in 2008, the number of active students was 30, of whom roughly half attended the lectures. From this group, 11 students participated in the self-assessment and diagnostic tests. A total of 23 students handed in the first assignment on time. The subsequent assignments were handed in on time by 15 and 16 students, respectively. Moreover, 8 out of the 11 students participated in the final written examination held in January 2009.

A. Self-Assessment Results

In the self-assessment test, the average test scores of the 11 respondents were above 5 for each of the featured subjects (see Table I for a summary of the results), with the ordinal scale of the test ranging from 1 (no knowledge) to 9 (extensive knowledge). Thus, in general, the students seem to perceive their own knowledge of the prerequisite subjects as above the level of an average understanding. According to the results, the students perceive themselves as having a good knowledge of basic mathematics and calculus, as well as of linear algebra. Not surprisingly, the students rate their knowledge of probability and statistics as just slightly above the average level of understanding.

B. Diagnostic Test Results

The average diagnostic test results, summarized in Table I, seem to correlate with the self-assessment test results in that most students have trouble answering basic statistics and probability questions, while the basic mathematics and calculus questions were answered correctly by the majority. The only significant difference between the two tests was that the majority of the students felt confident in their knowledge of linear algebra, but their diagnostic test results were rather poor. These results could of course be attributed to the fact that the linear algebra questions only cover a small part of the subject matter. Since the number of students participating in the study is low, the possibility should not be ruled out that the majority of students would fail even if they did have an average understanding of the subject matter. It is therefore obvious that just two questions for each subject are not enough to assess subject knowledge accurately. However, the rate of correct answers at least indicates the level at which the students will understand the concepts when they appear in the course literature.

C. Study Results

The study results for each student participating in the self-assessment and diagnostic tests can be viewed in Table II. This table actually provides data for the following metrics: self-assessment test score, diagnostic test score, completed assignments, and written examination test scores. Contrary to the design of Table I, the results provided in Table II are not averages but individual results for each participating student. These metrics have been normalized to simplify comparison

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

4   IEEE TRANSACTIONS ON EDUCATION

TABLE II
INDIVIDUAL RESULTS

Self assessment[1]	Diagnostic test[2]	Assignments[3]	Written examination[4]
0.84	0.20	0.67	N/A
0.80	0.90	1.00	0.72
0.78	0.50	0.67	0.42
0.78	0.40	0.67	N/A
0.78	0.60	1.00	0.64
0.76	0.70	1.00	0.42
0.73	0.40	0.67	0.64
0.73	0.20	0.67	0.50
0.67	0.20	1.00	0.64
0.62	0.50	1.00	0.53
0.58	0.10	0.67	N/A

[1] Normalized self-assessment test
[2] Diagnostic test
[3] Normalized score indicating the number of completed assignments (out of three) for all respondents
[4] Written examination scores

TABLE III
ACHIEVEMENT OF LEARNING OBJECTIVES

Learning Objective	Course Evaluation Test Score[1]
LO1	3.36 (0.89)
LO2	3.93 (1.03)
LO3	3.50 (0.91)
LO4	3.71 (0.88)
LO5	3.50 (1.35)

[1] Mean and standard deviation

V. ANALYSIS

The complex relationship between prior knowledge, prior experiences of learning and teaching, and hard work was touched upon earlier. More specifically, it seems that a suitable balance of these three factors is needed in order to achieve good study results. However, as has been pointed out, these good results might not coincide with a student's having obtained the necessary knowledge to understand the subject, and being able to use what has been learned in a practical situation.

A. Interviews

In order to get a more elaborate understanding of the relationship between prior knowledge and study results, two students were interviewed. The interviews were conducted after the course was completed and the students were graded, but before the students were aware of what grade they had received. From now on, these interviewed students are referred to as Respondent A and Respondent B. Respondent A emphasized the importance of having prior knowledge of artificial intelligence, back-propagated neural networks (a widely known machine learning technique), and calculus. The course literature [13] was described as being simple, practical, and not very mathematically intensive. However, one of the biggest hurdles in the course was reportedly the lectures on evaluation methods and evaluation metrics, along with the chapter on the same topic in the course literature. Respondent A suggested that a quite thorough knowledge of statistics would be needed in order to understand machine learning evaluation. Indeed, concepts such as probabilistic and rank-based evaluation metrics may not be included in many introductory statistics curricula, and hence it may therefore be difficult to relate to these concepts without prior knowledge of more advanced statistical theory.

Additionally, Respondent A felt that success in completing two of the mandatory hand-ins was heavily dependent on good programming skills. Thus, if the student did not have adequate programming skills, it would not matter much if he or she had sufficient knowledge or capabilities related to the machine learning topic associated with the hand-ins. It is possible that the notion of good skills within the area of programming may differ substantially between different groups of students, especially from different learning environments. As previously mentioned, the engineering subjects at BTH have traditionally emphasized the need for practical programming exercises throughout the Bachelor's-level education. Consequently, at the Master's level, the students are expected by many teachers to be fluent in languages such as Java and C++. However, over the last decade, BTH has noted a clear decrease in the programming skills of the Bachelor's students as well. This decrease may be a result of the lower number of applicants to engineering education. Compared to the more technology-intensive Bachelor's courses at BTH, such as Operating Systems, Real-Time Systems, and so forth, the machine learning programming assignments are simple and require only rudimentary skills in Java programming. The high number of failed assignments might, as earlier suspected, be attributed to poor understanding of how to use the machine learning theory that has been learned, in practice. However, as the discussion on

across metrics. It should be emphasized that the data is too limited, in terms of the number of respondents, to allow proper statistical tests. Instead, the results should be interpreted as an example of how to visualize and compare prior knowledge with study results. That being said, it might still be interesting to observe that there is a correlation between the self-assessment and diagnostic test results ($r = .60$), at least for the group of students studied (this was hinted at in the previous subsection). Again, this would indicate that a student who rates himself or herself as knowledgeable in terms of the subjects, considered here as important prior knowledge, will indeed do well in the diagnostic test and vice versa.

Moreover, it is also interesting to point out that, out of the top five students (according to their own self-estimates), only two succeeded in passing the written examination. As for the remaining three students, two did not participate in the written examination at all and one student participated but did not pass the examination (the threshold for passing the written examination is achieving by a 0.60 ratio of correct answers). In addition, the results in Table II also indicate that the diagnostic test is an even worse predictor of how well the student will perform during the written examination. It turns out that, from the top five students according to the diagnostic test, only one managed to pass the written examination.

D. Course Evaluation Results

Table III presents results for the achievement of learning objectives from the students' perspective. Since the course evaluation participation is anonymous, the results are reported as averages of the answers given by all participants. Judging by these results it is evident that many students felt that the evaluation and comparison of learning algorithms (using statistical tests and different evaluation metrics) were among the most difficult tasks, or at least these tasks represented concepts that were difficult to learn how to solve. In contrast, the students felt that the implementation of learning algorithms (using their programming skills) was an easier task. Additionally, it seems like many students also felt comfortable in comparing algorithms theoretically as opposed to comparing them empirically via statistical tests and experiments.

programming skills now suggests, it might just be that many students, although knowing how to solve the machine learning problem, are hindered by their inability to design and implement the code to achieve this.

This very notion is conveyed by Respondent B who reviews the third assignment, which consisted of implementing Q-learning; a reinforcement learning algorithm. The theoretical concept of Q-learning in itself was quite complex, Respondent B explained, but the greatest feat was to develop the application programmatically. Thus, instead of focusing on the course reference literature in which the Q-learning algorithm was covered [9], Respondent B searched the Internet and read different tutorials on how to implement the algorithm practically. When asked about the concept of evaluation of machine learning algorithms, Respondent B suggested that it was not the evaluation methods themselves that were difficult to comprehend, but rather, it was the statistical testing that was hard to understand.

B. Discussion

In the short term, it seems unlikely that external circumstances will change significantly, in that the number of applicants to computer science programs will suddenly increase or that their knowledge of mathematics, programming, and statistics will increase. Thus, it is arguably important to consider what can be done in terms of improving the content and presentation of the course. The results and interviews conducted within this study suggest that more focus needs to be put on getting students to understand the concept of statistical tests. As it happens, the experimental environment that is used in the course (Weka) automatically performs the t-test from within its Experimenter tool. It may be beneficial to spend more lecture time on manually calculating this and similar statistics, the rationale being that many of the Master's students do not seem to have any practical knowledge in this area.

However, an overly strong emphasis on statistics is not the solution. Opinions have already been voiced [12] that many of the current research papers put too strong a focus on obtaining statistically significant results and a comparably weak focus on explaining why certain results were achieved. It is a vital but difficult task to instead foster a scientific approach to developing critical and curious minds. It seems that both the positive aspects and possible issues in using statistical tests can be learned by having students actually perform the calculations themselves, and then reason out the results, on some problems that they can relate to. However, this analysis and reasoning needs to be properly explained, supervised, and encouraged by teachers.

In the long term, some of the common misconceptions about artificial intelligence and machine learning could perhaps be dispelled by simply introducing these subjects earlier than is commonly done today. It has even been argued that these subjects can be introduced as early as during secondary education [3]. At the university level, it may be necessary to face the fact that the average student of today has a different set of skills and knowledge than was the case perhaps just a decade ago. Therefore, the design and implementation of the Bachelor's and Master's programs need to be addressed in order to increase their flexibility so as to improve the learning of students with a range of prior knowledge and differences in learning strategies and study patterns. Some work has already been done in this area at many universities, for example, by decreasing the amount of higher level mathematics required and by repeating college-level mathematics during the introductory mathematics course at the university. With regard to implementation, it is arguably important to review the current curricula of education programs in the light of the changing skill levels of incoming

students to understand how to improve their learning by applying innovations such as changing the order in which courses are given and increasing connectivity between sequences of related courses.

VI. CONCLUSION

In this study, a class of students was given the opportunity to participate in a self-assessment test in which they were given the task of rating their own perceived level of prior knowledge of statistics, mathematics, and programming. After completing the test, the participants were given a diagnostic test, which included basic questions from these subjects. The relationship between the results of these two tests and the study results were investigated. It can be concluded that no obvious correlation exists between prior knowledge and study results, at least not for the group of students participating and if the prior knowledge is assessed through the featured self-assessment and diagnostic tests. However, self-assessment tests have been proven, by large-scale studies in the US, to accurately predict future study results in mathematics. In other words, the method of using a self-assessment test to predict study results seems suitable, but the low number of participants in conjunction with the self-assessment test used in this study may be too limited to achieve accurate predictions. Future work includes the development of a self-report, inspired by the Mathematics Science Inventory (MSI), but intended for Machine Learning students. This self-report will be evaluated empirically by studying a larger set of students.

REFERENCES

[1] D. Barber, "Master's level machine learning: Public perceptions," in *Proc. 2008 Workshop Teach. Mach. Learn.*, Saint-Etienne, France, 2008.

[2] B. Divjak and Z. Erjavec, "Enhancing mathematics for informatics and its correlation with student pass rates," *Math. Educ. Sci. Technol.*, vol. 39, pp. 23–33, 2007.

[3] R. Gavaldà, "Machine learning in secondary education?," in *Proc. 2008 Workshop Teach. Mach. Learn.*, Saint-Etienne, France, 2008.

[4] S. Lee, M. C. Harrison, G. Pell, and C. L. Robinson, "Predicting performance of first year engineering students and the importance of assessment tools therein," *Eng. Educ.*, vol. 3, no. 1, pp. 44–51, 2008.

[5] N. Entwistle, V. McCune, and J. Hounsell, "Approaches to studying and perceptions of university teaching-learning environments: Concepts, measures and preliminary findings," Enhancing Teaching-Learning Environments in Undergraduate Courses Project, Higher and Community Education, School of Education, University of Edinburgh, U.K., Tech. Rep., 2002.

[6] J. B. Biggs, *Teaching for Quality Learning at University*, 2nd ed. New York: Open Univ. Press/Soc. for Research into Higher Education, 2003.

[7] C. de la Higuera, "Theoretical computer science for machine learning," in *Proc. 2008 Workshop Teach. Mach. Learn.*, Saint-Etienne, France, 2008.

[8] A. Ambroladze and J. Shawe-Taylor, "Core skills curriculum for information technology," Curriculum Development Programme of the PASCAL Network of Excellence series, Tech. Rep., 2005.

[9] T. M. Mitchell, *Machine Learning*. New York: McGraw-Hill, 1997.

[10] E. Mjolsness and D. DeCoste, "Machine learning for science: State of the art and future prospects: A computer science odyssey," *Science*, vol. 293, pp. 2051–2055, 2008.

[11] P. Langley, "Machine learning: An editorial," *Mach. Learn.*, vol. 1, pp. 5–10, 1986.

[12] C. Drummond, "Machine learning as an experimental science (revisited)," in *Proc. 2006 AAAI Workshop Eval. Methods Mach. Learn.*, Boston, MA, 2006, pp. 1–5.

[13] I. H. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd ed. San Mateo, CA: Morgan Kaufmann, 2005.

[14] P. Ginns, M. Prosser, and S. Barrie, "Students' perceptions of teaching quality in higher education: The perspective of currently enrolled students," *Studies Higher Educ.*, vol. 32, no. 5, pp. 603–615, 2007.