Comparing the Perspectives of Students and Professionals on Machine Learning Pedagogy

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Suhayla Ahmed

Fall 2023

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Kathryn A. Neeley, Associate Professor of STS, Department of Engineering and Society

Introduction - Machine Learning Pedagogy at UVA

At UVA, there are various professors teaching the machine learning class, and the responsibility to teach the class is passed on between semesters. While there are differences between the teaching styles of the professors, they still follow the traditional lecture format. However, this has had varying levels of success, with many people enjoying the format (theCourseForum). Some students, however, did not feel like they learned what they had expected from the class, whether it was because of the highly theoretical lectures or the knowledge gaps that became evident in pop quizzes (Ahmed, 2023). For some, it is hard to explain or understand why this is the case and they're not able to know their own learning needs.

Nevertheless, it is evident that students and professors who teach the class both have their differences in perspective when it comes to the best teaching methods and what makes students successful. Even quantitative factors like grade averages often are not a factual determiner of these different attributes for a class. Knowing this, there is clearly a difference in understanding between those who are learning and those who have already learned and are now teaching machine learning.

After acknowledging that such differences exist, it is important to also acknowledge the possible consequences if these differences are not addressed. If catered computer education is not emphasized, a considerable portion of CS graduates will have gaps in knowledge that will compromise all industries in need of their expertise. If experts are not able to observe the impacts of their teaching styles on those being taught, it is also possible that less people will be attracted to the area of machine learning as it may unnecessarily gain a reputation for being notoriously difficult. While no area of computer science is without challenge, there should still be accessibility to knowledge when it can be achievable.

Approaching a resolution for this issue requires looking at both student testimonies and expert perspectives to see where differences and similarities lie. While experts carry knowledge that is far from negligible, they still lack the novice perspective of a student and hence should pay attention to feedback closely. Student testimonies can provide insight on whether specific instances of instruction are rewarding, whilst expert testimonies provide perspective on what a practitioner of machine learning looks like. Another important aspect to consider is the transition from a novice to an expert in a field, and how exactly a student can follow such progress for their success. In this paper, I argue that one can comprehend how experts can reframe their current thinking to draw upon their novice experiences and approach students' queries with more empathy and a more effective approach for bridging knowledge gaps.

Section I - Differences in Understanding

Near the end of my spring semester, I conducted informal interviews with three fellow students in my machine learning class about their experiences for my capstone project (Ahmed, 2023). Through various conversations I had had with them prior to the interviews, I had a contextual understanding of what their thoughts were and could tie a lot of their complaints to prior conversations with them. When viewing the main feedback from these reviews, I could pick up on some similarities between what everyone said. When asked to rate their preparation for quizzes on a scale of 1 to 5, interviewees rated an average of 2. While people liked the overall workload of the class and the transparency in the syllabus of what the course material looks like, they were unhappy with how mathematical the material was and felt a disconnect between the lecture content and assignments. The following figure compiles an (aggregated) version of my findings in these interviews.

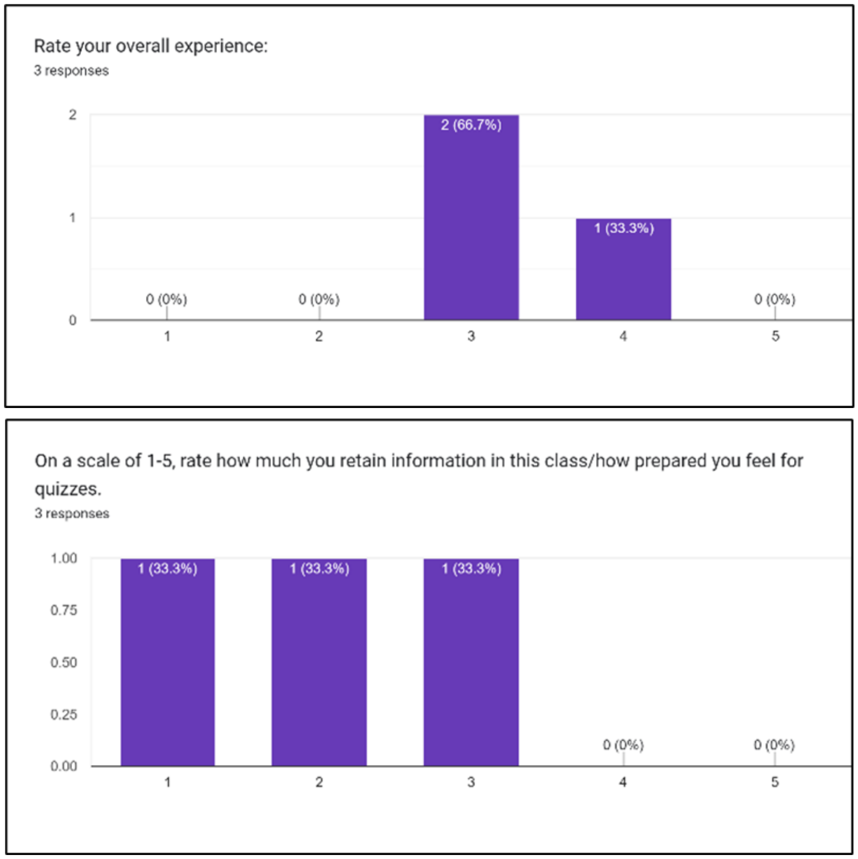

Figure 1: Findings From My Capstone Project About Student Experience in the Spring 2023 Semester of CS 4774 (Created by Author)

I also looked at theCourseForum for reviews on prior semesters. Although these were anonymous, the feedback looked specific enough to be able to connect to UVa's machine learning elective. There were a variety of professors appointed to teach the course, with each professor having a different approach and lecture style that would affect the way that they would be perceived by students. There were also differences in the teaching style of a COVID-affected online class versus a traditional in-person lecture. Compared to my in-person interviews from the spring semester, these reviews are much more diversified in terms of feedback. In my interviews with the students, all of them expressed facing difficulty with the course work. However, many students in theCourseForum gave the class good feedback, calling it "difficult but challenging" and praising the lecture style. There were still some complaints of not having the tools to succeed, i.e. lacking TA help or bearing the brunt of a highly mathematical approach to machine learning, but many reviews remained positive.

Knowing the contrast of reviews, it was important to think about the key differences between the spring semester specifically and other semesters the course was taught at UVa. Looking at the reviews for this semester on theCourseForum and at my reviews, there were many assignments that could not be given in time, either because they were not covered in lecture yet or because the faculty had not been able to create them in time for the class to complete it. As a result, the points were given back to the students, leaving students with more lost opportunities to exercise their knowledge. There was a different instructor for the spring semester from earlier ones, and his instruction style and assignments were received differently by students than by the professor of previous semesters.

The issue of a disconnect between professors and students' knowledge is nothing new. In nearly every class, higher education or otherwise, there is a student falling behind on course material due to a variety of reasons. According to a journal by the Department of Computer Science and Information Science at the Norwegian University of Science and Technology, there is an increasing demand for computer science professionals (Pappas et al., 2016). However, there was a shortage of CS professionals expected by the end of 2020 in Europe due to the high dropout rate (Pappas et al., 2016). In fact, it was estimated that a staggering 40 percent of students who pursued CS education would end up dropping out. There are various causes for this, but knowing that this is the case makes the emphasis on CS education all the more important as a more customized education would make students more enthusiastic in pursuing their profession. Pappas et al. hypothesize that a person's effort and performance determine their retention rate to a certain degree. At the same time, they also factor into account certain demographics. The figure below illustrates different hypothesized factors on a students' end affecting why they drop out, as their different demographics factor into their retention rate.

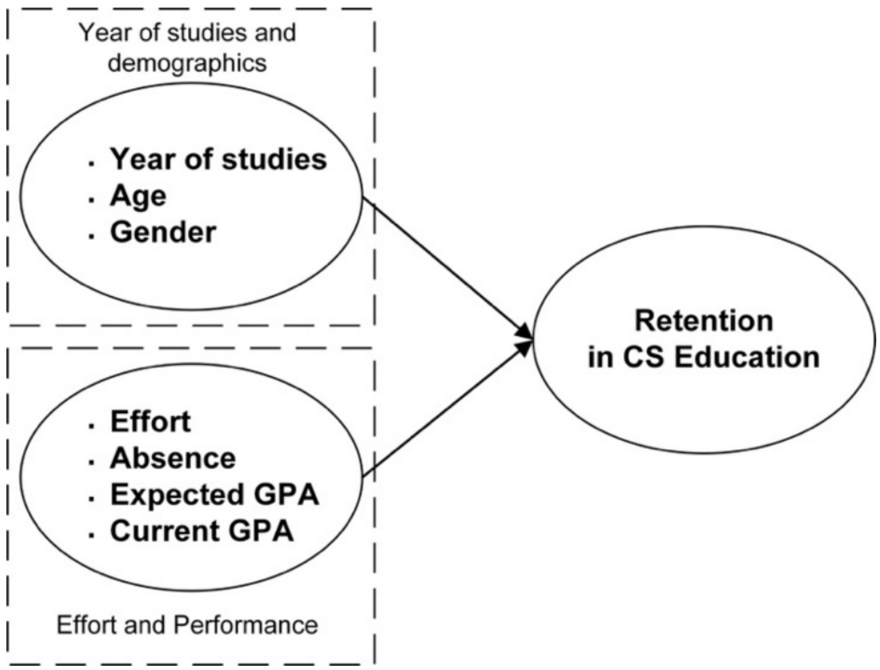

Figure 2: Conceptual Model for students' intention to complete their studies (retention) in CS education (Pappas et al., 2016)

Another important source to consider is the syllabi of machine learning elective classes across multiple universities. At the University of Virginia, the machine learning syllabus says the goal of the course is 'to understand basic concepts and algorithms in machine learning', with the subject matter desired to teach being rather mathematical aside from the introductory 'introduction to learning theory' topic listed (Ji, 2023). Delving outside of UVa, the syllabi of different machine learning classes were accessed to look at their different designs to teach students. Another prominent institution with a renowned CS department, New York University, had the syllabi for their machine learning class, with the course number being CS 6923 as it is a graduate school level class (NYU, 2021). At the beginning of the syllabus, the class is said to cover 'theoretical basics of a broad range of machine learning concepts and methods with practical applications to sample datasets via programming assignments'. The phrase 'principle models' is used more than once in the syllabus to describe the sort of mathematical concepts present within the subject. At Princeton University, there is an emphasis on exam-based learning in contrast to the project-based learning at UVA (Adams, 2023).

Another resource, which was not directly related to the pedagogical form of machine learning, was based on the comprehension of how machine learning systems can auto-generate explanations for people at the novice level of learning. From the American University of Beirut, it discusses methods to explain machine learning processes for non-experts (Diederich, 2020). According to the article, support vector machines, which are learning models in the area of machine learning, try to automatically generate explanations for certain behaviors in machine learning systems through more machine learning in a rather meta way. However, such analyses are not tailored to the needs of a novice user of said system. This leads to an impediment of the use of machine learning systems in general. The core idea of the article is extracting human-like media from the systems to improve comprehension. Such media can be movie clips or graphic illustrations.

Finally, resources were assessed on the process of developing from a novice to an expert in a topic. Patricia Benner, a registered nurse with a Ph. D., had a novice to expert theory that was in the realm of nursing. It was a way of understanding how a nurse, whether a novice themself or not, would acclimate to a particular skill or situation over time. The specific theory is applied to many situations in the nursing field, but is broad enough to the point where it can be applied to different situations. Additionally, Benner's theory takes inspiration from the Dreyfus Model of Skill Acquisition. The entire process is described as a circular rather than linear process since people learn many different skills in their lives, hence bringing them back to the novice stage. These specific resources also gave insight on which analytical approach to use when viewing the students' and experts' perspectives on machine learning.

Section II - Analysis of Differences Through a Novice to Expert Framework

As described in the previous section, the problem of finding a computer science class difficult is not an uncommon one. Perhaps the fact that it is so common makes an approach to trying to resolve this issue all the more daunting, as people are unsure of 'fixing' a system that seems to work well enough - there are still many CS graduates every year. However, to think about why sometimes a class fails a student in trying to gain more knowledge, one can approach an analysis of the two ends of a spectrum - the novice on one end, and the expert on the other.

In approaching an analysis of these perspectives, I will be using the analytical framework of Patricia Benner's Novice to Expert Theory (1982), which encompasses the process of skill acquisition. It is described to be applicable to 'a broad variety of situations, including nursing education, retention of graduate nurses, and nursing management and administration.' Benner, whose background is in nursing, provides a framework which is applicable to multiple disciplines and draws upon engineering and philosophy as sources of knowledge. Hence, it can be considered an interdisciplinary theory that has applications beyond its source. The overall process of skill acquisition is something human rather than tied to any profession.

The framework outlines five stages of proficiency, from the initial stage of knowledge acquisition to the 'final stage'. While it can be debated that the process from a novice to an expert is better outlined as a continuum rather than an arbitrary number of levels or stages, Benner's theory also argues that learners do not necessarily have to move through these stages in a systematic manner. In fact, some stages can be passed through multiple times as people gain more insight into new skills. The process is not a linear one, but rather a circular one. A summarized version of this model is shown in Figure 1. It shows the stages of novice and expert and the steps taken that take one from one end to the other.

Stage	**Definition**	**Knowledge acquisition**
Novice	The learner has had no previous experience	Teach simple concepts/attributes
Advanced beginner	The learner has enough real-world experience to understand themes in rules and guidelines.	Increase assistance and support
Competent	The learner has been on the job two or three years and is able work efficiently	Offer inservice education or opportunities
Proficient	The learner uses pieces of evidence (i.e. maxims) that provide directions to view a situation as a whole.	Use case studies to stimulate critical thinking
Expert	The learner grasps the situation and understands what needs to be accomplished beyond rules, guidelines, and maxims.	Provide opportunities for experts to share their skills and knowledge

Figure 3: Benner's Novice to Expert Theory in a table

In understanding how experts in machine learning came to be and also the level of understanding that new students begin with, I will be examining different pedagogical resources that detail the students' point of view and the experts' point of view. The student testimonies are not only indicative of the novice level, but of other levels above depending on how much information the students retained from the class. The differences and commonalities in student perspectives also provide a glance into how exactly they were successful or unsuccessful in transitioning through the different stages. This piece of information is fruitful in aiding instructors, as this perspective will help them make many more students successful in that regard. To get a better grasp on the students' side, I will assess student testimonies of this spring's machine learning class, which I was also a student in. Moreover, I will assess testimonies from previous semesters, both through online forums and interviews with other students. To assess the side of experts, I will read an article on experts' perspectives on machine learning and also assess the learning expectations for UVA's machine learning class. Finally, I will view sources on the transition of a novice to an expert and document general findings about this process. The following figure shows my plan of action.

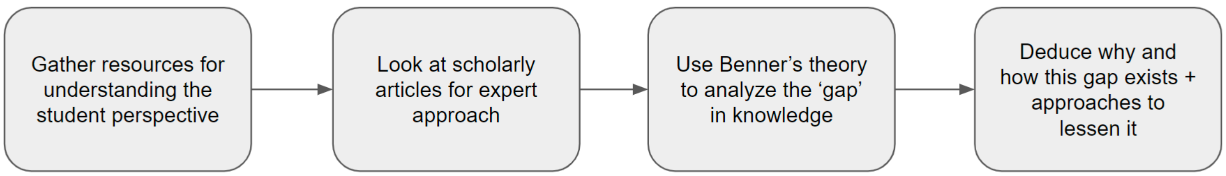

Figure 4: Flowchart of Step-by-Step process of the Methodology (Created by Author)

Section III - Results of Analyzing The Differences

Analyzing both novice and expert perspectives on machine learning provides valuable insight into the process of learning and how it can be more effective for those going into rising CS fields. First, I looked at the novice perspectives through interviews and theCourseForum. I look at expert perspectives through reviewing various syllabi for machine learning electives and various articles of expert perspectives on machine learning at a glance. I analyze the two extremes of novice to expert through Benner's theory.

To view experts' perspectives on machine learning as well, different sources were viewed about this discipline from an expert's lens. One resource discussed machine learning in radiation oncology (El Naqa, 2015). At a glance, machine learning was described as 'an evolving branch of computational algorithms that are designed to emulate human intelligence by learning from the surrounding environment'. As its premise, the goal of machine learning is to 'emulate the way that human beings learn to process sensory signals in order to accomplish a goal'. Keeping this in mind, it looks like expert perspectives about machine learning, even in an introductory sense, do not seem to attempt to make succinct explanations such that they're easily understood by people who are new to the field. This is perhaps because of the fact that science literature as a whole is not easily accessible by those who need it. In fact, scholarly articles are written in such a way that their alleged target audience may not even understand them themselves (Rosenberg et al., 2023; Hoeft, 2012). This begs to ask the question: "For whom is this literature for?" In fact, among non-humanities students, there is less inclination to be as well-versed in literature itself as opposed to disciplines more closely related to their major.

Education within the sciences is also vastly different from that of the humanities. As explained before, there is an emphasis on the mathematical aspect. Moreover, in the sciences there is a disconnection between what students learn in class and what they go on to do in industry. Even in the machine learning class at UVa and beyond, there was a common theme in the syllabi of the concepts learned being entirely mathematical with little connection to real life implications and real world impact. In a discipline like machine learning, one would not be shocked to see that there are vast real life implications outside of the classroom, with some begging for real ethical considerations.

To come close to an approach to resolving the problem, the gap between a novice and an expert must be accounted for. The first step to approaching a resolution is to understand that said gap is more of a continuum; one does not jump from a novice to an expert in time, but rather ebbs through the different phases outlined in Benner's theory, which itself may only encapsulate a select few phases of learning. When looking at the demographic of computer science students in UVA's machine learning class, there is a high chance that many of these students consider themselves experts, or at least proficient, in another area of computer science or another field entirely. The whole process of learning is a cycle that people go through for many different fields and subfields. Drawing upon one's past experiences learning something is important, as people's styles of learning are individual (Kolodner, 1983; C, 2023). However, there lies a deeper responsibility in professionals who are already teaching machine learning to understand strategies that make an entire audience better comprehend the intricacies of the elective. The next figure shows Benner's five stages again, but this time it shows what hinders the 'mobility' between the stages that can cause people to be stuck at earlier stages with an inability to progress.

Stage	**Hindrances to Knowledge Acquisition and Mobility**
Novice	The learner may feel constricted to abide by strict rules rather than move forward.
Advanced beginner	The learner is still constricted by rules and hence may feel an inability to use the real-world skills they have beyond a certain scope.
Competent	The learner starts to face difficulty abiding by strict rules and hence uses more real-world knowledge.
Proficient	The learner pushes out the rules in favor of their real-world experience.
Expert	Now as an expert, the learner is highly knowledgeable of the field. However, the learner may not be able to empathize with the earlier stages' inability to use real-world applications. For experts, rules are blurred rather than the set truth.

Figure 5: Benner's Novice to Expert Theory with Possible Hindrances to Knowledge Acquisition (Created by Author)

Novices in machine learning as well as in other fields are at a disadvantage for a variety of reasons, namely because they do not have the upperhand that experts have in terms of knowledge of a field and tools to grow their knowledge in said field. For example, in my interviews, when people were unsure of what they had gotten wrong in a certain homework assignment, they felt like the explanations they received were incomplete and did not augment their understanding of the topic. This especially is critical, as such an early adage cemented in people's brains (with truth to it, of course) is the idea that we all learn from our mistakes. Another example is asking for help at office hours, in a general sense. While oftentimes it is very helpful to go to office hours and get guidance from a teaching assistant on your homework, it also has its downsides. For example, teaching assistants may be too keen to give out homework solutions rather than guide students through problems themselves. This could stem from a variety of reasons, such as the idea that the assistants have not been able to grasp the concepts themselves well enough to be able to teach outside of the answer key. However, another possibility is that they simply believe giving the answers will be adequate, even though it does not serve to augment understanding. At the other extreme, however, some teaching assistants may explain in too convoluted a manner for students to understand.

The two extremes mentioned can be akin to the novice versus expert gap. A novice seems to follow things from whatever is given to them while a professional understands too much to the point where they can't effectively help those in a level below them. It is an interesting thought, as teaching assistants typically fall somewhere in the middle of Benner's five level model. They are certainly not novices, as they have been able to grasp the subject material for at least a semester. However, they may not be at the level of a professional either, as they may have just been introduced to the material through the class and may still have fallen victim to its shortcomings. These shortcomings, as mentioned before, can cause the class to be too abstract or too mathematical. While the abstractness is a strength in the fact that it is like the 'source' and is applicable to many scenarios, it can be a vice in the sense that it is simply a disconnect from what is going on in the perhaps dogmatically named 'real world'.

To make up for the gap in knowledge, experts and professionals have an opportunity to make their classes and expertise more accessible to those willing to learn. Course feedback is an excellent tool in gauging how to improve a class, and the most successful professors are often the ones who integrate feedback into their work (Watermark Insights, 2022). The feedback is also a way for students to self-reflect on how they could have better harnessed whatever tools professors gave them to learn.

Conclusion

In the deep investigation of the strengths and faults in UVa's machine learning class, many insights were gained into how instructors can rethink scientific pedagogy so that students are able to retain the most information possible and are well prepared for work in the future growing industry. The gap in knowledge can be rethought as a continuum, as new students will have different understandings from people who took the course, teaching assistants, and of course the professors. The experts of machine learning are able to define certain terms, but may not go into depth explaining their real-life applications and the like. There is a responsibility for instructors and professors to act in the interest of their students, as going against this would defeat the purpose of instructing a class. Viewing this gap of knowledge through the lens of Benner's Novice to Expert theory is fruitful. While it is logical to know that people are able to know much more when they are experts, it is an important distinction to make as people can make projects in machine learning completely incorrectly, which can be expensive in industry work.

Further looking through the lens of this theory, one can see that everyone goes through some process where they learn, but do not begin with an innate understanding of the field. This is especially true for machine learning since it is a developing field that has risen in the last century. While it is no secret to those working in technology that learning is a very tangible process of growth, the way to go about this process could still require some reframing of thinking. Both students and professors contain a reciprocal relationship in a class - not only do professors give feedback on student work, but students also must often give feedback on the professor's work, i.e. their teaching. Professors have another responsibility as the carriers of knowledge to better integrate feedback into their teaching so that the gap between them and students is less exacerbated. To expand upon this bridging of the gap further, more research could be done into the most effective strategies that cause students to truly learn in a class setting. This research can also be applied to fields beyond machine learning, and even areas beyond computer science; The issue of a gap between professors and students has always been interdisciplinary. All things considered, the best path forward beyond this paper is to listen more to what students have to say and enact that.

References

Adams, R. (2023). COS 324: Introduction to machine learning. Princeton University. https://www.cs.princeton.edu/courses/archive/fall18/cos324/

Ahmed, S. (2023, April 5). CS 4774 Learning Outcomes Interview 1. personal.

Ahmed, S. (2023, April 5). CS 4774 Learning Outcomes Interview 2. personal.

Ahmed, S. (2023, April 5). CS 4774 Learning Outcomes Interview 3. personal.

Benner, P. (1982). From Novice to Expert. Nursology. https://nursology.net/nurse-theories/from-novice-to-expert/

C, W. (2023, October 23). Novice expert scale - the theory + real life examples. YourDigitalAid. https://yourdigitalaid.com/novice-expert-scale-the-theory-real-life-examples/

Course syllabus - CS-GY 6923-Machine Learning. NYU Tandon School of Engineering. (2021). https://engineering.nyu.edu/sites/default/files/2021-06/Machine%20Learning%20Official%20Syllabus.pdf

El Naqa, I., & Murphy, M. J. (2015). What is machine learning? Machine Learning in Radiation Oncology, 3–11. https://doi.org/10.1007/978-3-319-18305-3_1

Hoeft, M. E. (2012). Why university students don’t read: What professors can do to increase compliance. International Journal for the Scholarship of Teaching and Learning, 6(2). https://doi.org/10.20429/ijsotl.2012.060212

Ji, Y. (2023). Machine Learning UVa CS 4774 (Spring 2023). Machine Learning. https://yangfengji.net/uva-ml-undergrad/

Kolodner, J. L. (1983). Towards an understanding of the role of experience in the evolution from novice to expert. International Journal of Man-Machine Studies, 19(5), 497–518. https://doi.org/10.1016/s0020-7373(83)80068-6

Pappas, I. O., Giannakos, M. N., & Jaccheri, L. (2016). Investigating factors influencing students’ intention to Dropout Computer Science Studies. Proceedings of the 2016 ACM Conference on Innovation and Technology in Computer Science Education. https://doi.org/10.1145/2899415.2899455

Rosenberg, A., Walker, J., Griffiths, S., & Jenkins, R. (2023). Plain language summaries: Enabling increased diversity, equity, inclusion and accessibility in scholarly publishing. Learned Publishing, 36(1), 109–118. https://doi.org/10.1002/leap.1524

theCourseForum. (2023). CS 4774. theCourseForum. https://thecourseforum.com/course/CS/4774/

Watermark Insights. (2022, March 16). The benefits of course evaluation in Higher Education. Watermark.

https://www.watermarkinsights.com/resources/blog/the-benefits-of-course-evaluation-in-higher-education