Are Interactive Visualizations in Machine Learning Education Helping Students?

Renea, Ilinca; Migut, Gosia; Krijthe, Jesse

DOI
10.1145/3724363.3729032

Licence
CC BY

Publication date
2025

Document Version
Final published version

Published in
ITiCSE 2025: Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 1

Citation (APA)
Renea, I., Migut, G., & Krijthe, J. (2025). Are Interactive Visualizations in Machine Learning Education Helping Students? In *ITiCSE 2025: Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 1* (pp. 2-8). ACM. https://doi.org/10.1145/3724363.3729032

Are Interactive Visualizations in Machine Learning Education Helping Students?

Ilinca Rențea
I.E.I.Rentea@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Gosia Migut
M.A.Migut@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Jesse Krijthe
J.H.Krijthe@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Abstract

With the fast integration of Machine Learning (ML) across industries, effective pedagogical strategies are essential for teaching this complex and evolving field. Machine Learning is now widely integrated into various university programs and introduced at earlier educational stages, including high school and secondary school. However, ML pedagogy lacks standardized teaching methods compared to other science-related subjects, which have established norms for topic introduction, teaching tools, and assessment methods. Inspired by other fields, this research explores the use of interactive visualizations in teaching ML topics, more specifically in teaching Gradient Descent (GD) and Principal Component Analysis (PCA). The target population consists of Computer Science and Engineering Bachelor students who have not yet followed any Machine Learning courses but have foundational knowledge in calculus, linear algebra, and statistics. The evaluation measures knowledge gained and student motivation, compared to a static version of the materials. Results show a significant positive effect in knowledge related to PCA with interactive visualizations, but no differences in knowledge gain for GD or in learning motivation for either topic. With these results, we contribute to the body of evidence-based teaching methods in Machine Learning and identify further research needed to generalize the effect of interactive visualizations as a teaching method for teaching ML basic concepts.

CCS Concepts

• **Social and professional topics** → **Computer science education**.

Keywords

machine learning, education, interactive visualizations, knowledge gain, motivation, controlled experiment

ACM Reference Format:
Ilinca Rențea, Gosia Migut, and Jesse Krijthe. 2025. Are Interactive Visualizations in Machine Learning Education Helping Students?. In *Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 1 (ITiCSE 2025), June 27-July 2, 2025, Nijmegen, Netherlands.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3724363.3729032

ACM ISBN 979-8-4007-1567-9/2025/06
https://doi.org/10.1145/3724363.3729032

1 Introduction

Machine learning (ML) is one of the most rapidly growing technical and research topics at the moment [20], having applications in various fields, such as healthcare [40], agriculture [23], education [21], and many more. Due to its popularity, much of the effort in the ML community is focused on advancing the field through application development and research. However, there has been much less emphasis on ML pedagogy and educational research, leading to a lack of evidence-based strategies to effectively teach a growing group of students these complex topics.

Learning objectives associated with this subject are **understanding** machine learning algorithms, **implementing** and **applying** them in specific use cases, **evaluating** the performance of such algorithms, and **analyzing** their performance and limitation, as suggested by [45]. However, these learning objectives can differ depending on the intended audience, since the knowledge can be applied in different manners and contexts [41]. More than that, machine learning topics are considered to be difficult because they rely on mathematics and abstract concepts [15]. In this context, interactive visualization may be a helpful tool to enhance understanding and engagement.

Firstly, teaching non-trivial mathematical concepts has been accompanied by using interactive visualizations. These are graphical representations that allow users to manipulate and explore the visualized information. The interactivity of visualizations enables the exploration of the underlying knowledge, resulting in a deeper understanding [8]. These tools were used in mathematical subjects such as calculus [16, 31, 38], geometry [24] and optimization problems [4]. Besides being effective in teaching mathematics, interactive visualizations have also been applied in teaching information retrieval [8], business analytics [37], and computer algorithms [9].

Secondly, students' engagement in the learning process is crucial for effective education [17]. In this context, interactive visualizations have been used to increase students' engagement, since they foster active learning, leading to a better understanding of complex topics [24, 26, 28]. Additionally, interactive tools encourage exploration and experimentation, leading to deeper cognitive processing and improved learning outcomes according to the ICAP framework [10]. However, despite promoting motivation and learning, interactive visualizations could introduce an additional cognitive load that is not related to the tasks that students need to perform [44].

Therefore, the research question this paper aims to explore and answer is the following:

> *How does the use of interactive visualizations in machine learning education affect students' knowledge gain and motivation related to machine learning topics?*

The initial hypothesis is that interactive visualizations will aid the students in understanding machine learning concepts. This hypothesis relies on the effectiveness of interactive visualizations in other similar subjects, such as mathematics and statistics, which are closely related to machine learning [42] , but also on their popularity in machine learning courses. Similarly, the second hypothesis is that interactive visualizations will positively influence the students' motivation related to machine learning topics.

We use two ML non-trivial topics to illustrate the influence of interactive visualizations on students' knowledge gain and motivation. The first topic is gradient descent, which was chosen because it was previously visualized [19, 47], but there is no evaluation of the performance of these visualizations. Whereas, PCA was chosen because it is considered to be a rather hard topic to teach to students, as suggested by Westfall [50].

The novelty of the current research does not lay in the use of interactive visualizations in the context of teaching machine learning, but rather in developing these visualizations and testing their effectiveness with students. To that end, the knowledge gained and motivation of students are measured when exposed to interactive visualizations, in comparison to when exposed to static visualizations.

The research paper will first touch upon the related work in section 2, then describe the methodology of developing the visualizations and evaluating them in section 3, followed by the results of the evaluation in section 4. Lastly, the discussion of the findings is included in section 5, followed by the conclusion and future work in section 6.

2 Related Work

Despite the importance of machine learning education, this research topic is only now beginning to take shape. One of the first steps taken in this direction is the agenda for future research developed by Shapiro and Fiebrink [41]. From the large number of research directions provided, the conceptualization and reasoning of students about machine learning algorithms and the parameters for these algorithms served as the starting point for our research. [41]. This point of the agenda highly focuses on exploring how students understand different concepts, which can be partially answered by exploring the usage of interactive visualizations.

Other researchers, like Skripchuk et al. [43], have focused on identifying common errors in open-ended ML projects. Their study primarily discusses code-related issues such as inadequate hyperparameter tuning and improper use of test data during model evaluation. However, the research focuses on errors in applying ML algorithms rather than students' understanding of the algorithms themselves.

The research field also features previous work focusing on the use of visualizations for teaching ML [11, 27, 48, 51]. The first example of relevant work done in this direction provides interactive visualizations to high school students [11]. In this context, visualizations are used for describing data collection, visualization, and processing, but also analysis, classification, and regression algorithms. All mentioned concepts are adapted to a daily life scenario, namely rain and weather prediction. However, this research does not evaluate the proposed method, leaving it unclear how effective or not this method would be when applied. Another relevant work for the current research is a convolutional neural network(CNN) visualizer [48]. The authors of this paper describe the process of developing the interactive visualization, while also including an observational study to evaluate the proposed solution. The evaluation is mainly focusing on the users' perception of the given visualization, rather than their true knowledge gain. The qualitative study's results included positive feedback from the participants, who found the method helpful in their study process, attractive, and easy to use. Another relevant example is the What-If Tool [51] which aims at helping users probe, visualize, and analyze different machine learning systems. This method is particularly relevant for the current research as it represents an interactive system that highly depends on visualizations since it focuses on minimal coding. The system's target user is represented by people working in different companies that interact with machine learning. However, the idea behind the system could also be used in an academic scenario. Interactive visualizations are also used in portraying fairness concepts related to machine learning algorithms and models, as studied by Mashhadi et al. [27]. Their research focuses on six open-source fairness tools, by conducting a qualitative review and four focus groups. Through their research, they show the importance of integrating interactive visualizations for teaching fairness in machine learning and artificial intelligence courses. One reported insight is that the previously mentioned What-If Tool [51] shows greater transparency due to its interactive design.

Interactive visualizations were also previously used in teaching other science, technology, engineering, and mathematics topics [29]. This teaching tool has been successfully applied in mathematics [1, 4, 16, 24, 31, 33, 38], algorithms education [2, 17, 36], statistics [22], information retrieval [8], and engineering education [39]. One key finding of previous research is that interactive visualizations aid in a better understanding of mathematical concepts [16, 24]. More than that, existing research shows that animations, non-interactive dynamic visualizations, might be helping students learn faster than static visualizations [9].

Lastly, existing research looked into the effect of different levels of interactivity on understanding and applying knowledge related to Signals and Systems in Electrical Engineering [35]. The results of the study showed that higher levels of interactivity did not have an effect on the learning outcome in understanding conceptual knowledge and applying procedural knowledge and had a negative effect on understanding procedural knowledge.

3 Methodology

3.1 Gradient Descent and PCA

To assess whether interactive visualizations positively impact the knowledge gained by students and their motivation, we first identify the key concepts of these two topics.

Gradient Descent is an optimization algorithm that is used ubiquitously in Machine Learning. This algorithm is used for searching through a large continuously parameterized hypothesis space when a performance measure can be differentiated with regards to the hypotheses [30]. It starts with a random value, updating the parameters of the hypotheses based on the gradient of the performance measure and the set learning rate.

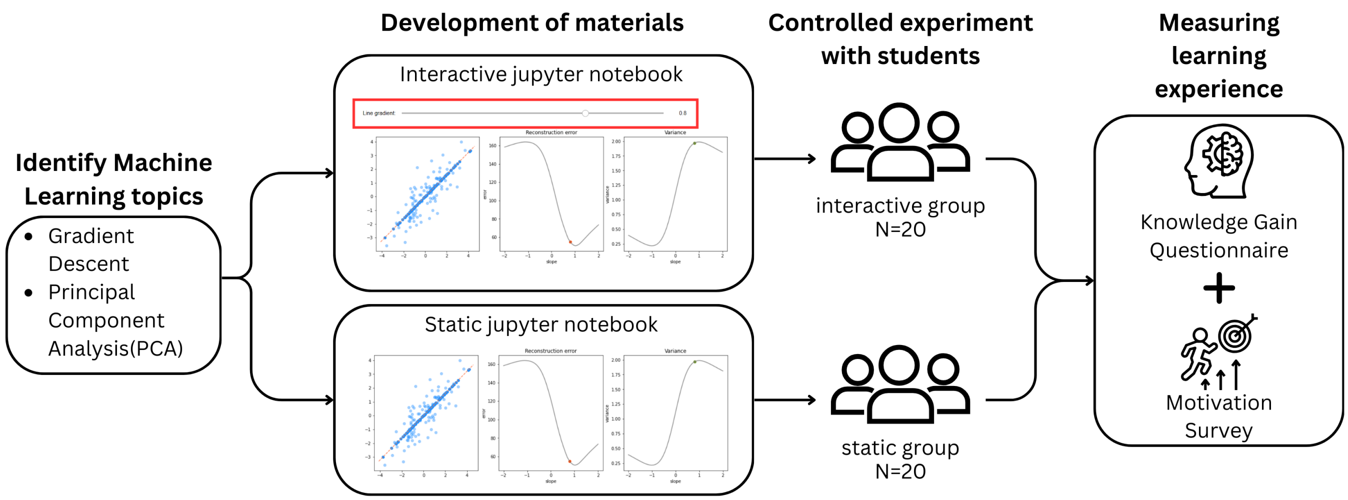

Figure 1: Detailed methodology behind research.

PCA is a dimensionality reduction algorithm commonly used in Machine Learning. This technique derives new variables that are uncorrelated linear combinations of the original variables of a given dataset, maximizing the retained variance of the dataset. The generated variables can be shown to be the eigenvectors of the covariance matrix of the dataset. These variables represent the principal components, ordered by decreasing eigenvalues, which are directly proportionate to the retained variance.

3.2 Interactive Visualizations Development

Starting from the concepts previously mentioned, two Jupyter Notebooks were created for the chosen topics. [1] The Gradient Descent notebook includes textual information, inspired by existing books and courses in Machine Learning [30, 32], and three interactive visualizations. The first visualization shows a one-variable scenario to introduce the students to the idea of gradient descent, and it allows them to set the starting point, learning rate, chosen function, and whether the gradient of the last update is shown graphically. The second visualization allows the students to manipulate the values of six parameters, aiming to help them understand the mechanics of gradient descent by manually performing coordinate descent, an algorithm based on the same principles as gradient descent. The third visualization displays a scenario with two variables, where the student can set the starting coordinates, learning rate, and epochs computed. The second and third visualizations are set within the scope of machine learning specifically since they display a regression scenario with a given dataset.

Similarly, the notebook introducing PCA includes textual information, inspired by books and courses in Machine Learning [6, 49], and three interactive visualizations. The first visualization displays a dataset, together with a line on which points are projected, but also the associated reconstruction error and variance of the performed transformation. The students in this case can choose the slope of the line used for the transformation. This visualization aims to portray the connection between reconstruction error and variance of transformations. The second visualization allows students to scale the 2 dimensions of the dataset, observing how the principal components are affected. The third visualization illustrates how data is transformed by the covariance matrix. The student can choose how many times the data is transformed and should note that the new points are along the principal component with the highest variance. This last visualization aims to make the covariance matrix a less abstract concept.

Lastly, cognitive load theory [46] was taken into account while developing the interactive visualizations. Since this teaching tool could sometimes introduce additional cognitive load [44], the visualizations' complexity was minimized while retaining the level of knowledge they delivered. This was done by only introducing necessary interactions and tackling subtopics through each visualization.

3.3 Experiment Design

The initial hypothesis needed to be tested using an experiment focusing on the knowledge gain and motivation of students. This section will focus on describing the methodology of the experiment.

Participants. The experiment recruited 40 first-year Computer Science and Engineering Bachelor students who did not follow any Machine Learning courses in their previous education and were the least likely to know gradient descent and PCA already. However, at the beginning of the experiment, participants were asked if they were already familiar with these two concepts to ensure that no participant had the knowledge before the experiment.

Procedure. The participants were divided into two groups: the static group (the control group, N=20), and the interactive group (the treatment group, N=20). The group assignment was done randomly. Each group received the same survey which included the pre-test and post-test questions, together with the Reduced version of IMMS (Instructional Materials Motivation Survey), for both topics. The procedure of giving a pre-test and a post-test to measure the knowledge gain is similar to the one applied by previous research [14, 34]. The pre-test and the post-test were built to match

[1] https://github.com/ieirentea/Interactive-Visualizations-ML

the instructional learning outcomes (ILOs) of the materials, according to the constructive alignment framework [5]. These outcomes were defined using Bloom's taxonomy [7], gradient descent being associated with understanding and applying the algorithm, and PCA with remembering and understanding it. The experiment was conducted in a quiet, distraction-free space, to ensure the optimum study environment for the participants.

Materials. For the experiment, the interactive group received written materials adapted from a bachelor-level Machine Learning course, accompanied by the interactive visualizations developed for this research. the static group received the same written materials, this time accompanied by static visualizations adapted from the developed interactive ones. The static visualizations were obtained by saving multiple instances of possible interactions with the visualizations. Due to the interactive nature of the second visualization of gradient descent, it was omitted from the static version of the materials.

Measurements. To measure the knowledge gained by the students, they were asked to fill in a pre-test, follow the given materials, and fill in a post-test. This procedure had the goal of measuring the knowledge gain of the students rather than the knowledge level, similar to the methodology proposed by [13, 18]. Both tests included the same questions, in increasing order of difficulty, as suggested by existing research [3]. The knowledge gain is measured using the following formula [12]:

$$g = \frac{post - pre}{1 - pre}, \tag{1}$$

where pre and $post$ are the scores of the knowledge tests. These scores were calculated based on the correctness of the answers, where each question received an equal weight for the final score. This relative measure of gain accounts for prior knowledge and enables fair comparisons.

After the post-test, students were given a questionnaire to measure their motivation regarding the received materials. The selected questionnaire is the Reduced Instructional Materials Motivation Survey [25], comprised of 12 questions that measure attention, relevance, confidence, and satisfaction associated with the materials. This survey was chosen due to its structured and validated development, together with its compatibility with the context of this experiment, namely a self-directed instructional setting.

Data Analysis. The collected data is analyzed using different statistical tests. The first step of the analysis is checking the assumptions of t-tests, namely normal distribution and equal variances. If the conditions failed, another appropriate statistical test is chosen (Mann-Whitney U-test or Welch's t-test).

Moreover, for knowledge gain, the pre-test scores were checked to not significantly differ between groups. This test ensures that the groups are similar in terms of pre-knowledge of ML.

4 Results

4.1 Knowledge gain

The first concept measured by this experiment is the knowledge gained by reading and interacting with the educational materials.

In figure 2, you can see the aggregated knowledge gain for each topic, comparing the two groups, where static is the control group and interactive is the treatment group.

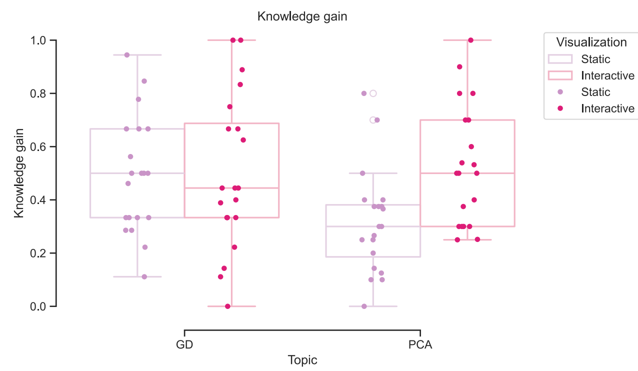

Figure 2: Knowledge gain results. The calculation of the scores is based on equation 1

Knowledge Gain gradient descent. Pre-test scores were non-normally distributed (Shapiro-Wilk: interactive W=0.5763, p<.001; static W=0.6637, p<.001). Therefore, a Mann-Whitney U-test was conducted and showed no significant difference between the pre-test scores of the two groups (U=190.5, p=.7639), namely the static (M=0.1028, SD=0.1759) and the interactive (M=0.1333, SD=0.2520) groups.

The normalized knowledge gain of the two groups was compared using a t-test due to the normally distributed data (Shapiro-Wilk: interactive W=0.9582, p=.5091; static W=0.9633, p=.6114) and equal variances (Levene's test: F=1.3963, p=.2447). The independent samples t-test showed no significant difference in knowledge gain between the interactive (M=0.5014, SD=0.2917) and static visualizations (M=0.4915, SD=0.2192) for the gradient descent topic (t(df)=0.1216, p=.9038).

Knowledge Gain PCA. Pre-test scores were non-normally distributed (Shapiro-Wilk: interactive W=0.5929, p<.001; static W=0.751, p<.001). Therefore, a Mann-Whitney U-test was conducted. The results of the test showed no significant difference (U=156, p=.1626) between the pre-test scores of the interactive (M=0.0532, SD=0.1003) and static (M=0.1050, SD=0.1276) visualizations groups regarding PCA.

The normalized knowledge gain of the two groups was compared using a t-test due to the normally distributed data (Shapiro-Wilk: interactive W=0.9246, p=.1215; static W=0.9331, p=.1769) and equal variances (Levene's test: F=0.7702, p=.3857). The independent samples t-test revealed a very significant difference in knowledge gain between interactive (M=0.5274, SD=0.2268) and static visualizations (M=0.3162, SD=0.1957) for the PCA topic (t(38)=3.1522, p=.0032).

4.2 Motivation

The second construct measured is the students' motivation regarding the instructional materials. The Reduced Instructional Materials Motivation Survey results were aggregated using the accompanying recommendations, namely calculating one score for the overall motivation, but also one score for each of the four constructs.

Table 1: Statistical Test Results for Knowledge Gain and Motivation

Variable	Group	Mean	SD	t/U	p-value
Knowledge Gain (GD)	Interactive Static	0.50 0.49	0.29 0.22	0.1216	.9038
Knowledge Gain (PCA)	Interactive Static	0.52 0.32	0.23 0.20	3.15	.0032
Motivation (GD)	Interactive Static	3.67 3.52	0.61 0.69	221.5	.569
Motivation (PCA)	Interactive Static	3.48 3.52	0.90 0.57	-0.193	.8485

The aggregated results for the gradient descent and PCA materials are displayed in figure 3, and 4 respectively. In both graphs, the overall motivation score is labeled 'motivation'.

Motivation Gradient Descent. The motivation scores regarding gradient descent materials were compared using a non-parametric Mann-Whitney U-test due to the non-normal distribution of the static group scores (Shapiro-Wilk: interactive W=0.9716, p=.7876; static W=0.8960, p=.0348). No significant difference was found between the interactive (M=3.6750, SD=0.6099) and the static (M=3.525, SD=0.6920) groups (U=221.5, p=.5690).

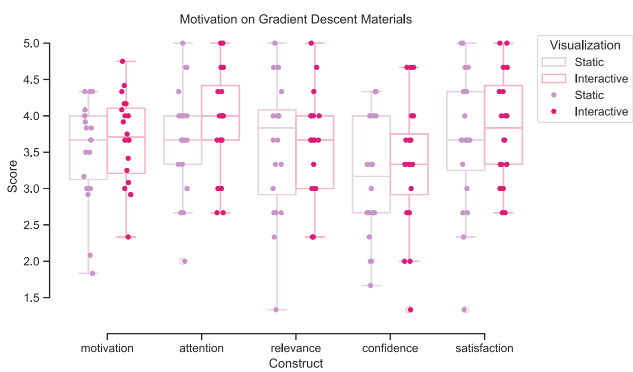

Figure 3: Motivation regarding gradient descent materials. Motivation is the overall aggregated RIMMS score and the rest are the 4 constructs measured by the survey, each aggregated separately.

Motivation PCA. A Welch's t-test was performed to compare the motivation regarding PCA materials due to unequal variances of the two groups (Levene's test: F=7.8823, p=.0078) and normally distributed data (Shapiro-Wilk: interactive W=0.9271, p=.1355; static W=0.9674, p=.7002). No significant difference in motivation regarding PCA was found between the interactive (M=3.4792, SD=0.8987) and static (M=3.525, SD=0.5698) groups (t(32.15)=-0.1926, p=.8485).

5 Discussion

5.1 Sample Size and Power Considerations

One of the key limitations of this study is the relatively small sample size (N=20 per group). This limited sample size may have reduced

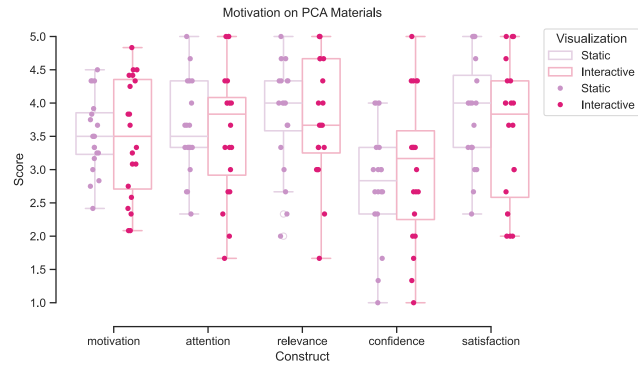

Figure 4: Motivation regarding PCA materials. Motivation is the overall aggregated RIMMS score and the rest are the 4 constructs measured by the survey, each aggregated separately.

the statistical power of the study, making it harder to detect smaller effects. For an effect to be statistically significant with this sample size, it would need to be relatively large. It's possible that smaller, more subtle effects of interactive visualizations on both knowledge gain and motivation were present but went undetected.

5.2 Knowledge gain

The results of the conducted study reveal no significant difference in knowledge gained by students between the interactive and static visualizations regarding gradient descent. This outcome indicates that the interactivity of the visualizations did not influence students' understanding of the presented materials.

On the other hand, the study revealed a very significant difference in knowledge gained by students regarding PCA. This result indicates that the added interactivity positively impacted the understanding of the presented materials, helping students better grasp the concepts related to PCA.

Therefore, the interactive visualizations had different effects on the two chosen topics. The difference between the outcomes could be a result of the different natures of the two topics. Gradient descent is a topic based on calculus concepts, closer related to previous knowledge of the students, while PCA is a topic based on linear algebra concepts, harder to understand and conceptualize, as mentioned by [50]. Based on the results, the interactive visualizations might help students better understand non-trivial and more abstract topics.

However, the difference could also originate from the design of the visualizations. In hindsight, it is worth considering whether both visualizations were equally well-constructed. Variations in clarity, interactivity, or complexity might have influenced the effectiveness of the materials.

The results of the study contradict the results of the existing research looking into the effect of different levels of interactivity [35]. However, that research focused on topics in Electrical Engineering, more specifically signal transformation, which involves manipulating concrete, real-world signals. This topic may be more accessible to the students compared to the ML concepts explored in this study, which are more abstract and rely on advanced mathematics, such as multivariate calculus and linear algebra.

Nevertheless, it is important to acknowledge that the observed differences in outcomes may not be only due to the chosen topics. The design and implementation of the interactive visualizations could have played a role in the results, as well as the prior knowledge and familiarity students had with each subject.

5.3 Motivation

The results of the conducted research reveal no significant difference in motivation levels regarding the constructed materials for gradient descent and PCA. Therefore, the introduction of interactive visualizations does not have a clear effect. One possibility is that the introduction of interactive visualizations does not have an effect on the motivation of the students. Another possibility is that the effect is too small to be observed using the current sample size, namely 20 students per group. However, it is important to mention that both versions of visualizations received relatively high motivation scores, having an average of around 3.5 (between moderately and mostly true), which might indicate little room for improvement.

The lack of a significant difference between interactive and static visualizations may be due to both types of visualizations having a similar impact on student motivation. However, it's important to note that motivation in this study was measured using the RIMMS survey, which gathers subjective responses from students. As a result, the findings reflect the students' personal opinions and perceptions of the instructional materials rather than objective measures. To increase the reliability of these findings, a larger sample size would be beneficial. A larger sample can help ensure that the results are more representative of the broader population and reduce the influence of potential biases, such as individual differences in how students interpret and respond to the survey.

6 Conclusion & Future Work

Future Work. The first limitation of the current research is the selected sample of the study. The experiment included 40 Computer Science and Engineering bachelor students. Expanding the number of participants and their diversity could strengthen the validity and applicability of the results in a more general context.

Another important limitation of the presented research study is the short-term nature of the controlled experiment. Due to time limitations, the controlled experiment fully focuses on short-term knowledge gain. However, in the future, it would be important to study the effects of interactive visualizations on long-term memory, since that is the main target of education.

More than that, for the current experiment, the setting was a controlled one, where the interventions were minimal and targeted on certain topics. However, the results of the experiment could be extended by conducting an in-the-wild experiment at one or multiple universities. This would imply inserting similar notebooks, maybe more than those presented in the current research, in Machine Learning courses, and observing their effect on students' learning experience, mainly through knowledge level and motivation.

Additionally, the current research only measures the effect of interactive visualizations for teaching gradient descent and PCA. These two topics are only a subset of the possible topics of machine learning. Therefore, further research is needed to explore the potential of these tools in the context of other topics.

Lastly, the focus of our experiment was on Computer Science Bachelor students. However, the effect of interactive visualizations should be also measured with students from other levels of education, such as pre-university, but also Master's students. More than that, the effectiveness of these tools might differ for non-Computer Science students, which motivates the need for this research direction as well.

Conclusion. This research looked into the effect of introducing interactive visualizations in teaching machine learning topics, namely gradient descent and PCA. Two sets of materials explaining the two concepts were created, one including static visualizations and one including interactive ones. The materials were used in a randomized controlled experiment with 40 Computer Science and Engineering Bachelor students. The experiment focused on measuring the knowledge gain and motivation related to each topic. The results of the study showed a significant difference in knowledge gained by students regarding PCA, interactive visualizations positively impacting the outcome. However, all other statistical tests returned no significant difference between the interactive and static visualizations.

Based on the results, interactive visualizations could be a useful tool in teaching certain ML topics. Due to no negative effect observed, such visualizations could be introduced in machine learning courses from computer science programs. However, the current research is only the first step towards validating the effect of using interactive visualizations in machine learning education. Additional research is still required to validate the effect of such visualizations regarding other topics commonly presented in machine learning and of other types of visualizations.

References

[1] Mojeed Kolawole Akinsola and IA Animasahun. 2007. The Effect of Simulation-Games Environment on Students Achievement in and Attitudes to Mathematics in Secondary Schools. *Online Submission* 6, 3 (2007).

[2] Nouf M Al-Barakati and Arwa Y Al-Aama. 2009. The effect of visualizing roles of variables on student performance in an introductory programming course. In *Proceedings of the 14th annual ACM SIGCSE conference on Innovation and technology in computer science education.* 228–232.

[3] Lina Anaya, Nagore Iriberri, Pedro Rey-Biel, and Gema Zamarro. 2022. Understanding performance in test taking: The role of question difficulty order. *Economics of Education Review* 90 (2022), 102293.

[4] Svetlana Asmuss, Natalja Budkina, et al. 2019. On usage of visualization tools in teaching mathematics at universities. *Engineering for Rural Development* 18 (2019), 1962–1969.

[5] John Biggs. 1996. Enhancing teaching through constructive alignment. *Higher education* 32, 3 (1996), 347–364.

[6] Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern recognition and machine learning.* Vol. 4. Springer.

[7] Benjamin S Bloom, Max D Engelhart, EJ Furst, Walker H Hill, and David R Krathwohl. 1956. Handbook I: cognitive domain. *New York: David McKay* (1956).

[8] Peter Brusilovsky, Jae-wook Ahn, and Edie Rasmussen. 2010. Teaching information retrieval with web-based interactive visualization. *Journal of Education for Library and Information Science* (2010), 187–200.

[9] Michael D Byrne, Richard Catrambone, and John T Stasko. 1999. Evaluating animations as student aids in learning computer algorithms. *Computers & education* 33, 4 (1999), 253–278.

[10] Michelene TH Chi and Ruth Wylie. 2014. The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational psychologist* 49, 4 (2014), 219–243.

[11] Siddharth Chittora and Anna Baynes. 2020. Interactive visualizations to introduce data science for high school students. In *Proceedings of the 21st Annual Conference on Information Technology Education.* 236–241.

[12] Elaine Christman, Paul Miller, and John Stewart. 2024. Beyond normalized gain: Improved comparison of physics educational outcomes. *Physical Review Physics Education Research* 20, 1 (2024), 010123.

[13] Michael Delucchi. 2014. Measuring student learning in social statistics: A pretest-posttest study of knowledge gain. *Teaching Sociology* 42, 3 (2014), 231–239.

[14] Timothy Ellis. 2004. Animating to build higher cognitive understanding: A model for studying multimedia effectiveness in education. *Journal of Engineering Education* 93, 1 (2004), 59–64.

[15] Christian Herta, Benjamin Voigt, Patrick Baumann, Klaus Strohmenger, Christoph Jansen, Oliver Fischer, Gefei Zhang, and Peter Hufnagel. 2019. Deep teaching: materials for teaching machine and deep learning. In *HEAD'19. 5th International Conference on Higher Education Advances*. Editorial Universitat Politècnica de València, 1153–1131.

[16] Andrea Hoffkamp. 2011. The use of interactive visualizations to foster the understanding of concepts of calculus: design principles and empirical results. *ZDM* 43 (2011), 359–372.

[17] Christopher D Hundhausen and Sarah A Douglas. 2002. Low-fidelity algorithm visualization. *Journal of Visual Languages & Computing* 13, 5 (2002), 449–470.

[18] Maria-Blanca Ibanez, Angela Di-Serio, and Carlos Delgado-Kloos. 2014. Gamification for engaging computer science students in learning activities: A case study. *IEEE Transactions on learning technologies* 7, 3 (2014), 291–301.

[19] Lili Jiang. 2020. A visual explanation of gradient descent methods (momentum, AdaGrad, RMSProp, adam). https://towardsdatascience.com/a-visual-explanation-of-gradient-descent-methods-momentum-adagrad-rmsprop-adam-f898b102325c

[20] Michael I Jordan and Tom M Mitchell. 2015. Machine learning: Trends, perspectives, and prospects. *Science* 349, 6245 (2015), 255–260.

[21] Danijel Kučak, Vedran Juričić, and Goran Đambić. 2018. MACHINE LEARNING IN EDUCATION-A SURVEY OF CURRENT RESEARCH TRENDS. *Annals of DAAAM & Proceedings* 29 (2018).

[22] Shailesh S Kulkarni, Bin Mai, S Yasaman Amirkiaee, and Hakan Tarakci. 2019. Dynamic interactive visualizations: Implications of seeing, doing, and playing for quantitative analysis pedagogy. *INFORMS Transactions on Education* 19, 3 (2019), 121–142.

[23] Konstantinos G Liakos, Patrizia Busato, Dimitrios Moshou, Simon Pearson, and Dionysis Bochtis. 2018. Machine learning in agriculture: A review. *Sensors* 18, 8 (2018), 2674.

[24] Hai-Ning Liang and Kamran Sedig. 2010. Can interactive visualization tools engage and support pre-university students in exploring non-trivial mathematical concepts? *Computers & Education* 54, 4 (2010), 972–991.

[25] Nicole Loorbach, Oscar Peters, Joyce Karreman, and Michaël Steehouder. 2015. Validation of the Instructional Materials Motivation Survey (IMMS) in a self-directed instructional setting aimed at working with technology. *British journal of educational technology* 46, 1 (2015), 204–218.

[26] Feiyu Lu, Difeng Yu, Hai-Ning Liang, Wenjun Chen, Konstantinos Papangelis, and Nazlena Mohamad Ali. 2018. Evaluating engagement level and analytical support of interactive visualizations in virtual reality environments. In *2018 IEEE international symposium on mixed and augmented reality (ISMAR)*. IEEE, 143–152.

[27] Afra Mashhadi, Annuska Zolyomi, and Jay Quedado. 2022. A Case Study of Integrating Fairness Visualization Tools in Machine Learning Education. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–7.

[28] Richard E Mayer. 2004. Should there be a three-strikes rule against pure discovery learning? *American psychologist* 59, 1 (2004), 14.

[29] Kevin W McElhaney, Hsin-Yi Chang, Jennifer L Chiu, and Marcia C Linn. 2015. Evidence for effective uses of dynamic visualisations in science curriculum materials. *Studies in Science Education* 51, 1 (2015), 49–85.

[30] Tom M Mitchell. 1997. Machine learning.

[31] Elena Nardi. 2014. Reflections on visualization in mathematics and in mathematics education. *Mathematics & mathematics education: searching for common ground* (2014), 193–220.

[32] Andrew Ng. 2000. CS229 Lecture notes. *CS229 Lecture notes* 1, 1 (2000), 1–3.

[33] Hitoshi Nishizawa, Takayoshi Yoshioka, Martti E Pesonen, and Antti Viholainen. 2012. Interactive worksheets for learning the connection between graphic and symbolic object representations. In *Proc. 17th Asian Technology Conference in Mathematics*.

[34] Barbara M Olds, Barbara M Moskal, and Ronald L Miller. 2005. Assessment in engineering education: Evolution, approaches and future collaborations. *Journal of Engineering Education* 94, 1 (2005), 13–25.

[35] Mrinal Patwardhan and Sahana Murthy. 2015. When does higher degree of interaction lead to higher learning in visualizations? Exploring the role of 'Interactivity Enriching Features'. *Computers & Education* 82 (2015), 292–305.

[36] Henry W Robbins, Samuel C Gutekunst, David B Shmoys, and David P Williamson. 2023. GILP: An Interactive Tool for Visualizing the Simplex Algorithm. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. 108–114.

[37] Dilal Saundage, Jacob L Cybulski, Susan Keller, and Lasitha Dharmasena. 2016. Teaching data analysis with interactive visual narratives. *Journal of Information Systems Education* 27, 4 (2016), 233–248.

[38] Kamran Sedig and Mark Sumner. 2006. Characterizing interaction with visual mathematical representations. *International Journal of Computers for Mathematical Learning* 11 (2006), 1–55.

[39] Levent Sevgi. 2006. Modeling and simulation concepts in engineering education: virtual tools. *Turkish Journal of Electrical Engineering and Computer Sciences* 14, 1 (2006), 113–127.

[40] K Shailaja, Banoth Seetharamulu, and MA Jabbar. 2018. Machine learning in healthcare: A review. In *2018 Second international conference on electronics, communication and aerospace technology (ICECA)*. IEEE, 910–914.

[41] R Benjamin Shapiro and Rebecca Fiebrink. 2019. Introduction to the special section: Launching an agenda for research on learning machine learning. 6 pages.

[42] R Benjamin Shapiro, Rebecca Fiebrink, and Peter Norvig. 2018. How machine learning impacts the undergraduate computing curriculum. *Commun. ACM* 61, 11 (2018), 27–29.

[43] James Skripchuk, Yang Shi, and Thomas Price. 2022. Identifying Common Errors in Open-Ended Machine Learning Projects. In *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education-Volume 1*. 216–222.

[44] Alexander Skulmowski and Kate Man Xu. 2022. Understanding cognitive load in digital and online learning: A new perspective on extraneous cognitive load. *Educational psychology review* 34, 1 (2022), 171–196.

[45] Elisabeth Sulmont, Elizabeth Patitsas, and Jeremy R Cooperstock. 2019. What is hard about teaching machine learning to non-majors? Insights from classifying instructors' learning goals. *ACM Transactions on Computing Education (TOCE)* 19, 4 (2019), 1–16.

[46] John Sweller. 2011. Cognitive load theory. In *Psychology of learning and motivation*. Vol. 55. Elsevier, 37–76.

[47] Teach LA Dev Team. [n. d.]. Gradient Descent Visualiser. https://uclaacm.github.io/gradient-descent-visualiser/

[48] Zijie J Wang, Robert Turko, Omar Shaikh, Haekyu Park, Nilaksh Das, Fred Hohman, Minsuk Kahng, and Duen Horng Polo Chau. 2020. CNN explainer: learning convolutional neural networks with interactive visualization. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1396–1406.

[49] Andrew R Webb, Keith D Copsey, and Gavin Cawley. 2011. *Statistical pattern recognition*. Vol. 2. Wiley Online Library.

[50] Peter H Westfall, Andrea L Arias, and Lawrence V Fulton. 2017. Teaching principal components using correlations. *Multivariate behavioral research* 52, 5 (2017), 648–660.

[51] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. 2019. The what-if tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 56–65.