

Adolescent and Adult Student Attitudes Towards Progress Visualizations

Onni Aarne*, Petrus Peltola*, Antti Leinonen, Juho Leinonen, and Arto Hellas

University of Helsinki
firstname.lastname@helsinki.fi

Abstract. Keeping students motivated for the duration of a course is easier said than done. Contextualizing student efforts with learning progress visualizations can help maintain engagement. However, progress can be visualized in many different ways. So far very little research has been done into which types of visualizations are most effective, and how different contexts affect the effectiveness of visualizations. We compare the effects of two different progress visualizations in an introductory programming course. Preliminary results show that older students prefer a visualization that emphasizes long-term progress, whereas younger students prefer a visualization that highlights progress within a single week. Additionally, students perform better and are more motivated when their visualization matches their age group’s preferred visualization. Possible explanations and implications are discussed.

Keywords: progress visualization, student engagement, growth mindset, motivation

1 Introduction

The effort that we expect from our students can be overwhelming to the extent that students temporarily lose track of their long-term goals and original motivations – seeing the forest from the trees can be hard. Progress visualizations can be a valuable tool for showing students the metaphorical forest, and helping them focus on the work that they are doing. Certain visualizations can improve students’ performance in a measured task; for example, a progress bar that measures a specific activity can increase students’ engagement with that activity [8]. These types of motivational tools can be important in labour-intensive courses, such as introductory programming courses, and can provide additional incentives for the students to complete the tasks they are given.

Learning progress visualization tools intended for students themselves have mostly been researched as part of broader attempts at gamifying education. However, this research rarely looks at specific elements of gamification with any rigor. In this work, we compare two types of progress visualizations that emphasize progress at different course granularities, and study whether specific visualizations are more suitable for specific course sub-populations. More specifically, we explore how adult and adolescent students respond to visualizations that highlight weekly progress and course-long progress.

* These authors contributed equally to this work.

This article is structured as follows. First, in Section 2 we provide an overview of students’ orientations and mindsets as well as discuss point and progress visualizations. Then, in Sections 3 and 4 we explain the details of the study and its results, respectively. In Section 5 we discuss the results and lastly, in Section 6 we conclude this work and outline possible future research.

2 Related work

In this section, we visit two relevant streams of research. First, we briefly discuss self-regulation in adults and adolescents. Then, we discuss previous work on motivational visualizations.

2.1 Students’ Mindsets as Factors of Success

This stream of research aligns with human self-regulatory processes, which broadly encompasses the ability of an individual to control and evaluate their behavior during learning. Caprara et al. [3] studied how changes in self-regulatory efficacy affects students’ grades. They found that as the efficacy declined, the students’ grades lowered. This in turn led to a higher chance of dropping out from studies. Their dataset consisted of 412 students, and the study was conducted as a longitudinal study over a period of ten years – the participants were followed between the ages of 12 and 22.

In her research, Carol Dweck observed that individuals who believe that intelligence is fixed, or that success is based on innate ability, are more likely to stop working during setbacks or when struggling than those who believe that intelligence can grow through effort [4]. Dweck popularized this belief using the term “mindset”. Later, different types of interventions that emphasize growth of knowledge over fixed knowledge have been developed.

For example, Rourke et al. [12] developed two versions of a math video game called “Fractions” to study the effects of growth mindset. Their experimental version of the game promoted students’ learning with commentary that praised their intelligence, whereas the control group got more neutral feedback praising their mathematical skills. They found that the experimental group were more persistent as they played for a longer period of time and completed more levels than the control group. They also started to display behavior coherent with a growth mindset by using strategies they learned from the game. Another result of the study was that lower performing students were more encouraged to keep playing in the experimental version than in the control version.

The goal is to provide the students visualizations that best support self-regulation. This in turn would lead the students to the correct mindset which prevents quitting.

Steinberg et al. [15] found that young adults’ future orientations and delay discounting continued to develop from childhood all the way into their mid-twenties. The research compared 10 to 30 year old individuals and indicated, when compared to their counterparts, that younger individuals are more likely to choose a smaller reward sooner than a larger reward later. This supports our study design.

2.2 Point and Progress Visualizations

Overall, point and progress visualizations fall under multiple umbrellas. They are an integral part of gamification research that seeks to bring game-like elements into educational contexts as well as more generic educational principles of transparency where students' progress should be made visible.

Leppänen et al. [8] studied the effect of progress visualization on student engagement. They found that, at least initially, a simple progress bar significantly increased the quantity of completed exercises. Their system displayed the bar at the bottom of the screen, and the bar became full when the user had answered three consecutive multiple choice questions correctly. However, the arbitrary nature of this goal meant that their bar simply created an illusion of progress – which was sufficient for increasing students' effort – rather than depicting anything more meaningful.

Loboda et al. [9] developed and studied a form of open learner modelling which is comparable to our progress visualizations, though more clearly split up by subject. However, in their study, the use of the system was entirely optional so any apparent effect might simply be due to better students being more likely to use the system.

Another study of open learner modelling was conducted by Bull et al. [2]. They used multiple visualization types for the same student, and the students could choose which visualizations were visible to them. The available visualizations included skillometers, tables, stars, smileys, gauges, histograms, networks, radar plots, word clouds and treemaps. These visualizations had information about students' knowledge and skills on the subject. They found that students prefer simpler visualizations such as skillometers and stars over complex visualizations such as treemaps and word clouds.

Santos et al. [14] studied similar use of visualizations. Their research had a few differences from ours. Their study was iterative. In the first phase they collected feedback from teachers. The teachers valued simplicity and intuitivity, which is in line with Bull et al.'s results [2]. After modifications between evaluations, they provided the visualizations to students. They found that their demographic, which consisted of 18 to 20 year old students, valued the visualizations. Another difference to our research was that they showed the students their peers' progress, which the students found positive.

Visualizations have also been used in education to present students' progress to teachers [5, 7]. Park and Yo [13] studied different usages of learning analytics dashboards for students, teachers or both. Based on the information they collected from previous studies of dashboards, they developed their own dashboard for both students and teachers, with the focus on students' self-regulation in the variables shown in the dashboard. They developed their dashboard iteratively similarly to Santos et al. [14] by interviewing students and improving the dashboard based on the feedback. Their dashboard also allowed students to see each other's progress.

Gamification in education [6, 10, 11] utilizes visualization of students' progress. For example Holman et al.'s [6] learning management system GradeCraft shows students their course scores as a progress bar. Additionally, they can get badges

as they learn new concepts which are also shown in the progress bar. Grade-Craft shows useful visualizations for teachers as well. They can see which badges have been completed, and which are in progress. They also see how the students predicted their success during the course and if their progress corresponds the prediction.

3 Methodology

In this section, we first explain the context of our study in Section 3.1. Then, we describe the survey the students were given and the students' answers, i.e. the data for our study in Sections 3.2 and 3.3. Lastly, our research questions are outlined in Section 3.4.

3.1 Context

The data for this study comes from a CS1 course series split into two seven-week courses conducted in the spring of 2016 at the University of Helsinki. The course series is an introduction to object oriented programming with Java with content similar to most CS1 courses globally. Since most CS majors take these courses in the fall, our sample mostly consists of non-CS majors.

The course uses a blended online material with interspersed theory sections and assignment sections. For the purposes of this study, a visualization component was added to the online material.

In the study, the students were randomly divided into two groups, which were shown separate visualizations of their progress. One group was shown a series of bars that corresponded to the percentage of completed exercises for each week of the course, and the other group was shown a line plot that emphasized the growth of students' skills throughout the course. The visualization aggregated students' scores over the course, showing a steady growth as students' progressed throughout the course.

From this point onward, we refer to the visualization with the weekly bars as the *Weekly visualization*, and the visualization with an emphasis on course-long growth and effort as the *Growth visualization*. Two examples of the Weekly visualization are shown in Figure 1, while Figure 2 shows two examples of the Growth visualization.

The visualization was implemented as a floating element that was constantly available in the lower right corner of the course material. As the students worked on the course, they were able to hide the visualization temporarily. Regardless of whether the student had previously hidden the visualization, the visualization was visible to the student when the material page was opened again.

3.2 Survey

At the start of the final exam of the course, the students were asked to fill out a survey about the course. As an incentive, they were offered an extra point toward the exam score. All of the survey questions relevant to our study were Likert items scaled from 1–7, where a low number indicated agreement and vice versa. A zero option was also included if the students felt the item was not applicable to




Fig. 1. Examples of the Weekly visualization, where the percentage of completed assignments for each week of the course are shown as a bar chart. The student on the left did fairly well, whereas the student on the right could not quite keep up but did not give up.




Fig. 2. Examples of the Growth visualization where aggregate assignment score is displayed on the x axis and time is shown on the y-axis. On the left is a perfect score, and on the right the student slowly fell behind and gave up.

them. The items for this study included statements about whether the students saw the progress visualization, whether they liked the visualization they saw and whether they felt motivated by it. The distributions of the responses to the questions are shown in Figures 3–6.

3.3 Data

Our study uses data from two subsequent programming courses that are typically taken together. In our dataset each student in a course makes up a single data point. Therefore students who took both courses contributed two points of data. Our survey from the course exams provided data from 89 and 83 students for each course respectively. For the study, we only included students who agreed that they saw the visualization, i.e. whose responses were between 1 and 3 on the 7-point Likert scale. This narrowed our dataset down to a total of 118 data points. We had to exclude 27 data points due to missing information on the participants’ ages, leaving us with a final dataset of 91 data points.

We divided the data into two groups: adolescents and adults. We consider adults to be students born before the year 1995, i.e. students who turned 22 the year of the study. This cutoff was chosen since it formed quite even groups and kept the age of the adolescents reasonably low. A later cutoff year would have raised the adolescent group’s mean age to over 20. With our chosen cutoff the adolescents were 18–20 year olds, while the adults’ age ranged from 21 to 34 years. The medians, means and standard deviations of the ages for the different groups are reported in Table 1.

There were 34 adolescents and 57 adults. Of all the individual students, 17 only took the first course, 20 only took the second course and 27 took both. For the students who took both courses, the visualization was switched between courses, which means that 27 students reported their experiences with both visualizations. The number of students in different age groups for the two visualizations can be seen in Table 1.

Table 1. The group sizes for the different age and visualization groups, along with the medians, means and standard deviations of the groups’ ages.

Age group	Weekly Vis.	Growth Vis.	Median Age	Mean Age	σ of Age
≥ 21	27	30	23	24.11	3.64
< 21	11	23	20	19.56	0.56
Total	38	53	21	22.41	3.64

3.4 Research Questions

We have the following two research questions for this work:

- RQ1. Does the type of visualization affect student scores?
- RQ2. Do participants of different ages respond differently to the different progress visualizations?

With the first research question, we are interested in studying whether the two visualizations affect students' scores differently, that is, whether the type of visualization matters. With the second research question, we examine whether a student's age affects how the student is affected by the visualization.

4 Results

To get a general overview of the results and distributions of the relevant survey responses, we visualize them in Section 4.1. We then analyze the results in greater detail, and seek to answer the research questions in Sections 4.2–4.4.

To answer research question 1, we normalized the total assignment scores for each student and calculated the means and standard deviations for each age group by visualization, reported in Table 4.

For research question 2 we calculated means and standard deviations of the responses to the survey questions “Did you like the visualization?” and “Did the visualization motivate you to complete more assignments?”, which are reported in Tables 2 and 3 respectively.

4.1 Questionnaire responses

The distributions of the questionnaire responses for each age and visualization group are reported above in Figures 3–6.

4.2 Preference

The different age groups seemed to prefer different visualizations. When asked if they liked the visualization, the mean response level of adolescents who saw the



Fig. 3. The questionnaire response distribution from the adults who saw the Weekly visualization.




Fig. 4. The questionnaire response distribution from the adults who saw the Growth visualization.



Fig. 5. The questionnaire response distribution from the adolescents who saw the Weekly visualization.




Fig. 6. The questionnaire response distribution from the adolescents who saw the Growth visualization.

Weekly visualization was 1.55, whereas the adolescents’ opinion on the Growth visualization was only 2.00. This indicates that the adolescents collectively preferred the Weekly visualization to the Growth visualization. Conversely, the mean response levels of adults were 2.22 and 1.65 for the Weekly and Growth visualizations respectively, i.e. the adults liked the Growth visualization more than the Weekly visualization.

There is very little difference between the standard deviations for the Growth visualization. For the Weekly visualization the standard deviations are significantly different between the two age groups: 0.69 for the adolescents and 1.72 for the adults. These values are reported in Table 2.

4.3 Motivating

The age groups were also more motivated by their preferred visualization. Adults reported being more motivated by the Growth visualization. When asked to report if the visualization affected their motivation to complete assignments, the mean response levels were 3.04 for the Weekly visualization and 2.19 for the

Table 2. Likert scale responses to the question “Did you like the visualization?”. The responses were between 1 (completely agree) and 7 (completely disagree).

Age	Weekly Vis.		Growth Vis.	
	Mean	σ	Mean	σ
≥ 21	2.22	1.72	1.65	1.28
< 21	1.55	0.69	2.00	1.38

Growth visualization. Similar to the preference question, the effect was reversed for adolescents. Their mean response levels were 2.27 for the Weekly visualization and 2.96 for the Growth visualization.

The differences in standard deviations follow the same trend that appears in the preference subsection, only to a lesser extent. The Growth visualization’s standard deviations are not significantly different, whereas the Weekly visualization’s standard deviations are 2.01 for the adults and 1.79 for the adolescents, i.e. there is a slight difference. Although the standard deviations of the Weekly visualization group differ for both of the questions, the difference is significantly higher for the preference question. These values are reported in Table 3.

Table 3. Likert scale responses to the question “Did the visualization motivate you to complete more exercises?”. The options were between 1 (completely agree) and 7 (completely disagree).

Age	Weekly Vis.		Growth Vis.	
	Mean	σ	Mean	σ
≥ 21	3.04	2.01	2.19	1.64
< 21	2.27	1.79	2.96	1.75

4.4 Points

The normalized assignment scores for the two different visualizations and age groups are reported in Table 4. We see a similar effect, albeit weaker, as with motivation and liking: younger students complete slightly more assignments when shown the Weekly visualization (86% completed vs 80% completed) and older students complete marginally more assignments when shown the Growth visualization (84% completed vs 81% completed).

Table 4. Average scores for the different age and treatment groups normalized to be between 0 and 1.

Age	Weekly Vis.		Growth Vis.	
	Mean	σ	Mean	σ
≥ 21	0.80	0.13	0.84	0.10
< 21	0.86	0.08	0.81	0.09

5 Discussion

Our results indicate that the long-term focused visualization received more positive responses from adult students and was associated with higher scores com-

pared to the alternative. Similarly, the visualization which emphasizes weekly progress received more positive responses and resulted in higher scores in the younger group, compared to the alternative.

Therefore, the answer to our first research question, “*Does the type of visualization affect student scores?*”, is affirmative. Moreover, as younger and older students preferred different visualizations, a more specific answer that also answers our second research question, “*Do participants of different ages respond differently to the different progress visualizations?*”, is the following: “*Depending on the age of the student, different progress visualizations affect students’ scores differently*”. However, the differences between the visualizations were subtle. This is especially the case when comparing our results to Leppänen et al. [8], who found a much more intense effect. A number of relevant differences may explain this. Their study used simple multiple choice questions, and kept the progress bar always visible – the progress bar also indicated when the student had “completed” a specific task, even though in reality their completion was arbitrary. In our study, the progress was slower as the assignments were more complex, and the visualization was only visible in the web-based course material and not in programming environment where students worked on most of the course tasks.

It is possible that a greater effect could be achieved if the visualization was shown after every completed assignment, also within the programming environment in which the students completed their course work.

We found that older students like the Growth visualization that emphasizes long-term gains more than the Weekly visualization that highlights short-term progress. These results are in line with those reported by Steinberg et al. [15] who found that students’ future orientation continues to develop until their mid-twenties. This suggests that there might be a benefit to adjusting motivational tools according to the students’ ages, even among legal adults.

Generally, students both liked and were motivated by the two visualizations. For all groups, the mean response level was under 4 on a scale from 1 (completely agree) to 7 (completely disagree). Thus, our results indicate that students like both visualizations, and the differences in our results are only in the magnitude of the liking. This is in line with previous work by Bull et al. [2], who found that students enjoy simple visualizations, as both of the visualizations in this study can be considered simple. On the other hand, it is possible that students who did not like the visualization turned the visualization off, and therefore answered that they did not see the visualization – such students were excluded from this study.

Another reason for the younger students thinking more favourably of the Weekly visualization might be that it is more familiar to them. Progress bars are commonly used for depicting various elements such as skill proficiency and enemy health in video games, and presumably younger students spend more time on them on average compared to older students.

In a broader context, our results indicate that in order to increase students’ learning, it might be more effective to personalize visualizations depending on

students' demographic factors such as their age. Even subtle improvements in students' success could be beneficial in e.g. programming courses that suffer from notoriously poor pass-rates [1].

5.1 Limitations

In this subsection, we acknowledge the limitations of our work. As our study is based on a small population with rather small differences in the outcomes, none of the results are statistically significant, despite the visual analysis that indicates otherwise. Therefore, this study should be treated more as qualitative first insight into this topic rather than a quantitative examination of the effect of different visualizations on learning. Comparisons between adult and adolescent learners are always difficult as the age is rarely the only difference; older students are often at very different points in their lives. Due to our method of splitting the students into only two age groups of under 21-year-olds and at least 21-year-olds, the younger group is much more homogeneous than the older one. For example, there is likely a larger difference between a 22-year-old student and a 32-year-old student compared to a 22-year-old student and a 20-year-old student, but in our study the 22-year-old would be in the same group with the 32-year-old student.

Our study also lacked a control group of students who saw no visualization at all. As such, we can not say how much the existence of a visualization affected students performance over a baseline – we can only compare the effect between the two visualizations. However, because most of the students we studied said they liked and were motivated by the visualization they saw, we can assume that the visualizations did not have a negative impact on the students' scores. That is, we may assume that a control group would have performed more poorly.

6 Conclusions

In this article we studied how different age groups responded to being shown different visualizations of their course progress – overall, students liked the visualizations and they can be a useful tool for educators in keeping students motivated and engaged. However, our preliminary results suggest not all such visualizations are made equal. Out of our two progress visualizations, the one that emphasized long term growth was better liked and promoted better scores among the older age group, whereas the younger age group performed better with a bar plot visualization that emphasized each week's progress individually. The differences were rather subtle, however, and were not statistically significant.

In the future, we would like to further test our findings on a course with more students. It would also be valuable to have a control group who were not shown any visualization to study how these visualizations compare to not having a visualization at all. In addition, we want to study how increases in the number of completed exercises translates to actual acquired skills, for example whether there are differences in exam scores between the different visualization groups.

References

1. J. Bennedsen and M. E. Caspersen. Failure rates in introductory programming. *ACM SIGCSE Bulletin*, 39(2):32–36, 2007.
2. S. Bull, B. Ginon, C. Boscolo, and M. Johnson. Introduction of learning visualisations and metacognitive support in a persuadable open learner model. In *Proceedings of the sixth international conference on learning analytics & knowledge*, pages 30–39. ACM, 2016.
3. G. V. Caprara, R. Fida, M. Vecchione, G. Del Bove, G. M. Vecchio, C. Barbaranelli, and A. Bandura. Longitudinal analysis of the role of perceived self-efficacy for self-regulated learning in academic continuance and achievement. *Journal of educational psychology*, 100(3):525, 2008.
4. C. S. Dweck. *Mindset: The new psychology of success*. Random House Incorporated, 2006.
5. A. L. Dyckhoff, D. Zielke, M. Bültmann, M. A. Chatti, and U. Schroeder. Design and implementation of a learning analytics toolkit for teachers. *Journal of Educational Technology & Society*, 15(3):58, 2012.
6. C. Holman, S. Aguilar, and B. Fishman. Gradecraft: What can we learn from a game-inspired learning management system? In *Proceedings of the third international conference on learning analytics and knowledge*, pages 260–264. ACM, 2013.
7. M. E. Jacovina, E. L. Snow, L. K. Allen, R. D. Roscoe, J. L. Weston, J. Dai, and D. S. McNamara. How to visualize success: Presenting complex data in a writing strategy tutor. In *EDM*, pages 594–595, 2015.
8. L. Leppänen, L. Vapaakallio, and A. Vihavainen. Illusion of progress is moar addictive than cat pictures. In *Proceedings of the Third (2016) ACM Conference on Learning @ Scale, L@S '16*, pages 133–136, New York, NY, USA, 2016. ACM.
9. T. D. Loboda, J. Guerra, R. Hosseini, and P. Brusilovsky. Mastery grids: An open-source social educational progress visualization. In *Proceedings of the 2014 Conference on Innovation & Technology in Computer Science Education, ITiCSE '14*, pages 357–357, New York, NY, USA, 2014. ACM.
10. B. B. Morrison and B. DiSalvo. Khan academy gamifies computer science. In *Proceedings of the 45th ACM technical symposium on Computer science education*, pages 39–44. ACM, 2014.
11. S. O'Donovan, J. Gain, and P. Marais. A case study in the gamification of a university-level games development course. In *Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference*, pages 242–251. ACM, 2013.
12. E. O'Rourke, K. Haimovitz, C. Ballweber, C. Dweck, and Z. Popović. Brain points: a growth mindset incentive structure boosts persistence in an educational game. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 3339–3348. ACM, 2014.
13. Y. Park and I.-H. Jo. Development of the learning analytics dashboard to support students' learning performance. *J. UCS*, 21(1):110–133, 2015.
14. J. L. Santos, S. Govaerts, K. Verbert, and E. Duval. Goal-oriented visualizations of activity tracking: a case study with engineering students. In *Proceedings of the 2nd international conference on learning analytics and knowledge*, pages 143–152. ACM, 2012.
15. L. Steinberg, S. Graham, L. O'Brien, J. Woolard, E. Cauffman, and M. Banich. Age differences in future orientation and delay discounting. *Child Development*, 80(1):28–44, 2009.