Dijkstra's Algorithm: Comparing PS-I and I-PS in an Online Learning Experiment

[redacted]

Professors:
Pierre Dillenbourg
Patrick Jermann

June 12, 2025

Contents

1 Introduction **1**

2 Learning Goals **1**

3 Lesson Design and Activities **1**
3.1 Overall design 2
3.2 Instruction Design 2
3.3 Problem Solving (PS) Activity Design 3
3.4 Pre-test and Post-test Questions 4

4 Experimental Design **4**
4.1 Research Question 4
4.2 Variables 4

5 Implementation **6**
5.1 Demographics and Feedback Surveys 6
5.2 Pre-test and Post-test 6
5.3 Instructions 7
5.4 Problem Solving Activity 7

6 Data Analysis **7**
6.1 Participants Data 7
6.2 Results 7
6.2.1 Exploratory Analysis 7
6.2.2 Control Variables 9
6.2.3 Power Analysis 9
6.2.4 One-way ANOVA 9
6.2.5 Moderation Effects 10
6.2.6 Mediation effects 12
6.2.7 Moderated Mediation effects 13

7 Discussion and Conclusion **14**

A Appendix **16**
A.1 Lesson design 16
A.2 Implementation 17
A.2.1 Surveys 17
A.2.2 Pre-test and Post-test Questions 18
A.2.3 PS Activity 18
A.3 Data Analysis 21
A.3.1 Mediation Analysis 21

1 Introduction

In this project, we compare the effectiveness of two instructional approaches—Problem Solving before Instruction (PS-I) and Instruction before Problem Solving (I-PS)—for teaching Dijkstra.

Topic Dijkstra's algorithm is one of the most well-known and widely used shortest path algorithms in a graph with non-negative edge weights. For example, if the nodes in the graph represent cities and the edge weights represent distances between cities connected by direct roads, Dijkstra's algorithm can be used to determine the shortest route from a *source* city to a specific *destination* city.

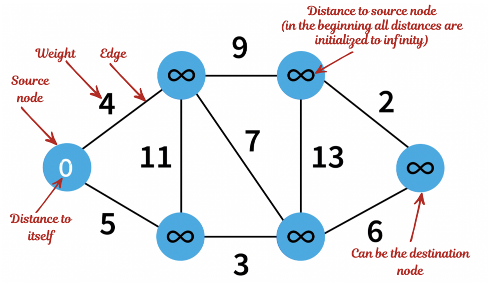

Figure 1: Example of a Graph at the Beginning of Dijkstra's Algorithm

Target Audience and Pre-requisites The lessons are aimed at young adults aged 18 to 30 with a scientific background. The participants are expected to be currently undertaking studies or to have completed their studies in the past 3 years. They should have little to no prior knowledge of shortest path algorithms, particularly Dijkstra's algorithm. They are, however, expected to have a basic understanding of graph theory, shortest path intuition, as well as the ability to engage in problem solving and critical thinking required to grasp these concepts and participate in the activity.

2 Learning Goals

The aim of this lesson is to help students understand Dijkstra's Algorithm, recognize its usefulness, and apply it to real-life problems. We have the following three learning goals:

- **LG1**: Describe how Dijkstra's Algorithm works (detailed steps) and its purpose.
- **LG2**: Compute the shortest path using Dijkstra's Algorithm.
- **LG3**: Apply Dijkstra's Algorithm to real-world problems by extracting and translating relevant information into a graph.

3 Lesson Design and Activities

This section outlines the design of our PS-I and I-PS experiments, detailing the lesson structure, activities, and their alignment with participants' prior knowledge and learning goals. We also discuss the theoretical foundations underlying the design setup and present the assessment questions.

3.1 Overall design

The entire experiment is designed to take 45 minutes to 1 hour to complete. Figure 2 shows the flow of the PS-I and I-PS experiments.

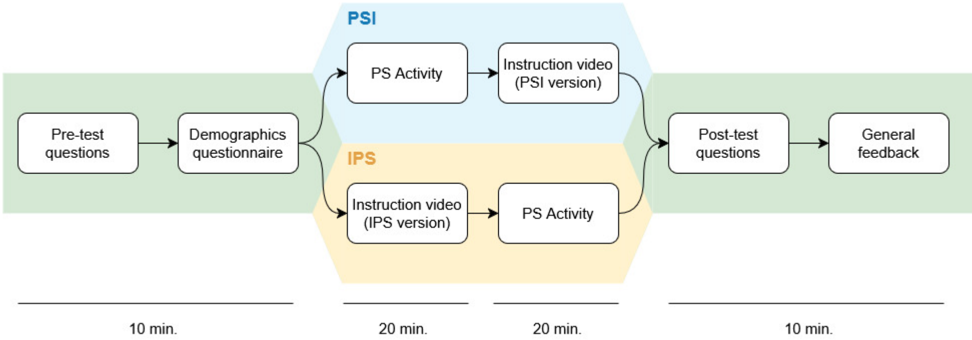

Figure 2: I-PS vs. PS-I Experiment Steps with Time Limits

3.2 Instruction Design

To accommodate the different needs of the PS-I and I-PS instructional sequences, we created two separate videos for the experiment. In the **PS-I condition**, the instruction video follows the problem-solving activity and explicitly solves the same problem as introduced in the PS activity, presented in the Appendix A.2.3. This approach enables a smooth transition from the activity to the instruction phase by avoiding **Split Attention Effect**. It also reduces potential frustration in participants who were unable to solve the activity and are curious to know the answer. In the **I-PS condition**, the instruction video precedes the PS activity and uses a different problem, while preserving the same conceptual structure of a shortest path navigation task. The example problem is illustrated in Figure 10 of Appendix A.1. This structure encourages learners to apply their knowledge about Dijkstra's Algorithm to a new context (i.e. train transportation systems) different from the one seen in the instruction phase. Both versions of the lesson follow the same structure, with each component aligned to one or more learning goals:

- **Introduce the purpose of Dijkstra's Algorithm and its real-life applications** (addresses **LG1**): This section helps students understand the importance of Dijkstra's shortest path algorithm by knowing its applications (e.g., Geographical Positioning System (GPS)).
- **Present two graph representations—adjacency matrix and adjacency list** (supports **LG1**): Revision of graph representations knowledge is essential to grasping how the algorithm operates.
- **Demonstrate step-by-step how to compute the shortest path using Dijkstra's Algorithm** (addresses both **LG2** and **LG3**): In both videos, we extract graph information from a table (PS-I) or a picture (I-PS) and apply the algorithm to solve it.
- **Briefly explain the time complexity of the algorithm** (completes **LG1**): Though not assessed, this provides a more complete understanding of the algorithm's properties.

To **reduce extrinsic cognitive load**, we made these design choices:

1. Videos emphasize visual clarity, through consistent color-coding and animations to represent recurring actions, such as node discovery and adjacency matrix updates.

2. We chunk the lesson content into several components, spending proper time per component.
3. We avoid the use of complex terminology or unexplained jargon.

To **maximize germane cognitive load**, we followed these practices:

1. We start by demonstrating different real-life applications for Dijkstra. This provides meaningful context, which increases motivation in participants and helps with schema formation.
2. We show Dijkstra's Algorithm steps via a worked example rather than just listing them.
3. We briefly review the prerequisites for our lesson, including various graph representations.

3.3 Problem Solving (PS) Activity Design

Our PS activity is a shortest path problem. Although we unified this activity for our two experiments (PS-I and I-PS), we designed it with the PS-I mode in mind. In other words, we avoided overguiding participants to foster **discovery learning**, in line with **constructivism**, by presenting the problem without explicitly framing it as a shortest path problem. Instead, we presented it as a real-life scenario, which serves three key goals:

- It makes the problem more intuitive and accessible, **reducing extrinsic cognitive load**.
- It enables learners to approach the problem by drawing on prior real-life experiences with transportation and route planning (i.e. **assimilation**).
- It helps **activate the prior knowledge** needed for the instruction phase - specifically, shortest path intuition, by connecting it to transportation routes, and algorithmic thinking.

The problem is designed so that the only way to reach the destination on time is by following the shortest path. Nonetheless, the participants are free to answer any way they want, so they might transform the problem into a graph representation or come up with their own algorithm, making the activity **well-suited** to the participants' **prerequisite knowledge**.
We present the information of city departures, city destinations, and travel times in a table (Figure 9 in Appendix A.1). To promote **productive failure** and **desirable difficulty without adding extrinsic load**, we designed the train table information to be challenging enough (e.g. no direct or single-hop path from source to destination), yet we limited the number of cities to only 10. Moreover, we follow two key design choices to promote **germane cognitive load** by encouraging meaningful problem-solving strategies, thus supporting **schema construction**. First, the problem is structured such that solutions following no algorithmic reasoning or any scientific intuition (e.g. pruning irrelevant cities) are significantly more time- and effort-intensive. Second, to prevent participants from relying on their prior geographical knowledge (e.g., knowing Lausanne is closer to Paris than New York), we used unfamiliar city names and unrealistic travel times to encourage focus on algorithmic thinking. Accordingly, learners are prompted to develop new ways to find the solution (i.e. **accommodation**).

We encouraged participants to extensively document their problem-solving process. Writing down their thoughts promotes participants to develop an awareness of their thought processes and refine their strategies to successfully discover the shortest path (i.e. **meta-cognition**), all the while helping us conduct a deeper analysis by tracing participants' solutions.
For PS-I, the activity allows learners to engage in deep analysis (**analysis** stage in **Bloom's taxonomy**). The insights gained from such an analysis will lay the foundation for introducing Dijkstra's algorithm in the instruction phase. For I-PS, it provides a good opportunity to apply

what they have just learned and ensure that they have understood the lesson (**application** stage in **Bloom's taxonomy**). Finally, the whole activity was designed to be completed within 20 minutes, with participants explicitly instructed not to exceed this limit. This time constraint helps keep the total experiment under one hour, maintaining participant engagement and reducing the dropout rate.

3.4 Pre-test and Post-test Questions

The content of both tests is identical and is composed of six questions requiring different levels of knowledge and collectively measuring all learning goals. The full list of the assessment questions can be found in Appendix A.2.2.

- **Question 1** evaluates general knowledge about shortest path algorithms by asking participants to identify real-world applications of Dijkstra's Algorithm.
- **Question 2** assesses specific procedural understanding of Dijkstra's Algorithm, focusing on the order in which nodes are processed.
- **Question 3** asks participants to find the shortest path in a directed weighted graph.
- **Questions 4 and 5** require detailed application of intermediate steps of Dijkstra's Algorithm to determine tentative distances and the next node to process.
- **Question 6** presents a short scenario with text-based route descriptions and travel times. Participants must determine the shortest travel time from a source to a destination.

4 Experimental Design

This section presents our main and sub-research questions and explains how we measure the different variables in our experimental design. These variables include independent, dependent, moderation, and mediation variables. We discuss how our experimental design, including the described variables, can answer our research questions.

4.1 Research Question

This study aims to answer this research question: **"Is there a statistically significant difference in relative learning gains between participants taught Dijkstra's shortest path algorithm using the PS-I versus I-PS instructional approaches?"**

Additionally, this study explores the following sub-research questions:

- Do individual difference factors (academic background, confidence level, academic level, parental education, and athletic experience) moderate the relationship between instructional approach and relative learning gains?
- Do cognitive and motivational processes during learning (problem-solving approach type, visualization strategy, and interest level) mediate the relationship between instructional approach and relative learning gains?

4.2 Variables

Independent Variable: The instructional approach, with participants assigned to PS-I or I-PS using a semi-random approach with constraints. Specifically, our goal was to balance the two

groups with respect to gender and academic background to avoid confounding the main analysis.
Dependent Variable: The relative learning gain, calculated as

$$\frac{\text{Post-test score} - \text{Pre-test score}}{\text{Max score} - \text{Pre-test score}}$$

This metric was chosen because it accounts for differences in participants' baseline knowledge. To calculate the relative learning gains, both groups completed the same pre-test and post-test questions designed to cover all learning objectives. (The full list of questions can be found in Appendix A.2.2.)

- **Questions 1 and 2** target **LG1** by focusing on understanding how Dijkstra's Algorithm works and its purpose in real-life applications.
- **Questions 3, 4, and 5** support **LG2** by focusing on computing the shortest path using Dijkstra's Algorithm.
- **Question 6** is aligned with **LG3** by applying Dijkstra's Algorithm to real-world problems by modeling them as graphs.

Control Variables: The time limit for completing each phase of the activity, as well as the core lesson content, were held constant across both groups. Academic background and level were originally intended to be controlled through selective recruitment. However, due to recruitment challenges, these variables were instead measured through a questionnaire at the start of the experiment and will be statistically accounted for in the analysis. Gender, however, has been balanced across groups.

Moderation Variables: Our experiment explores potential moderators to better understand the conditions under which or populations for whom the observed effect occurs. These variables were collected through direct questions in our pre-experiment demographic questionnaire:

1. **Confidence level** before starting the experiment: Students with higher baseline confidence may benefit more from PS-I's challenging approach, while less confident students might prefer I-PS's scaffolded structure.
2. **Academic background**: Students from computer science or mathematics backgrounds may benefit more from the PS-I approach compared to I-PS, as their stronger foundation in algorithmic thinking and graph theory allows them to more effectively engage with the problem-solving activity before receiving explicit instruction.
3. **Academic level**: Students at a higher academic level may benefit more from PS-I approach compared to students at lower academic levels due to higher confidence and more experience with problem solving.
4. **Parents' education level**: This may moderate the effect due to different levels of exposure to various teaching methods at home.
5. **Sports level**: Athletic experience may moderate the effectiveness of PS-I versus I-PS, as athletes might have more competence in problem-solving approaches. Studies have shown physical activity to be positively correlated with cognitive performance [1, 2]

Mediation Variables: Our experiment also explores potential mediators to better understand the mechanisms through which the approach influences learning gain.

1. **PS activity solution methodology** (Algorithmic versus Non-algorithmic & Graphical versus Non-graphical): We hypothesize that the participants' solution techniques can impact their depth of engagement with the activity. In particular, we hypothesize that techniques that are more graphical and algorithmic lead to greater engagement and, consequently, higher learning gains.
2. **Interest level**: Interest level is estimated indirectly via time spent on the assessments. Students with high prior knowledge (pre-test scores ≥ 7) are excluded, as they may spend less time. We acknowledge that time spent alone may not be the optimal proxy for interest, and in-person experiments would allow better assessment through facial expressions, and other body language, which is not feasible online. We hypothesize that the PS-I approach increases students' intrinsic interest more than the I-PS approach, which in turn leads to higher engagement and improved learning outcomes.

5 Implementation

The experiment was implemented on Moodle, in the form of a multi-section course. The goal was to allow the participants to complete the experiment at any time, without requiring our assistance and without having to meet the other participants.
Moodle pages can be accessed here for I-PS and here for PS-I. Participants were sent only the page of the version that was assigned to them.

5.1 Demographics and Feedback Surveys

The demographics and feedback surveys were implemented as Moodle quizzes. Our primary goal was to minimize the time and effort required from participants, while still allowing them to provide additional detail when desired. The **demographics questionnaire** was the first element participants completed. It collected basic background information, including gender, age range, level of education, field of study, parents' highest education level, and frequency of physical activity. The **feedback survey** was carried out at the end of the experiment. It gathered information about participants' perceived comfort in answering the assessment questions both before and after the instruction, their confidence level at the start of the experiment, and their perceived learning gains from the lesson.
This structure ensured that we captured relevant participant context and self-reported learning experiences with minimal burden on the participants. Figure 11 of Appendix A.2.1 displays a few examples.

5.2 Pre-test and Post-test

We implemented both the pre-test and post-test as Moodle quizzes. The format included multiple-choice questions (Questions 1 and 2), a drag-and-drop activity (Question 3), and short text answers (Questions 4, 5, and 6). Each question type was selected to align with the nature of the skill or knowledge being assessed. We attributed a maximum of 3 points to each learning goal, for a total score of 9 points. When a learning goal was assessed by multiple questions, the points were distributed equally among those questions. Given the 10-minute time limit, we emphasized that mistakes and uncertainty were acceptable to reduce pressure and discourage seeking external help. We also explicitly instructed participants not to use any external sources when answering. The "I don't know" option was included in MCQs to prevent guessing and ensure the reliability of measured learning gains.

5.3 Instructions

We recorded two videos using slides and voiceover. The video presented for I-PS can be found here, while the PS-I one is accessible here. We inserted the links in the Lesson section of the respective Moodle experiments.

5.4 Problem Solving Activity

We implemented the PS activity as an Assignment on Moodle. As mentioned above, the participants were strongly encouraged to join a picture of their scrap paper or tablet, to help us understand their thought process. Appendix A.2.3 discloses the exact handout that was given to the participants and shows visuals of the submission boxes on Moodle, plus an example of a participant submission.

6 Data Analysis

In this section, we describe our participants' demographics and analyze the data to answer our research questions (§ 4.1). Additionally, we use statistical tests and visuals to guide our analysis and draw insights. For this entire analysis, we use a confidence interval of 95%, i.e., the p-value should lie below a threshold $\alpha = 0.05$ for us to reject the null hypothesis. The code, dataset, and full analysis are accessible via this repository.

6.1 Participants Data

A total of 21 participants were recruited through our personal connections. However, one participant was removed from the analysis after achieving a perfect score in both the pre- and post-tests. The following analysis includes the remaining 20 participants. Participants were assigned to the instructional approach PS-I (n=10) or I-PS (n=10) using a semi-random approach, as explained in Section 4.2. Figure 3 shows basic statistics of our participants. Non-engineering fields include Economics (n=1), Environmental Laws and International Relations (n=1) and Medicine (n=1). Other engineering fields include Engineering and Entrepreneurship (n=1), Environmental Science (n=2), Life Sciences engineering (n=2), Mechanical Engineering (n=1), Micro-engineering (n=2), Neuro-X (n=1) and Physics (n=1).

6.2 Results

As explained in 6.2.3, a power analysis suggests that the small sample size impacts the significance of our measured statistics, so we conduct the analysis while reporting the relevant F and p-values even for high p-values. However, we try to draw interesting insights depending on the relative magnitude of F and p-values across different mediation and moderation effects.

6.2.1 Exploratory Analysis

To start our analysis, we conducted an exploratory analysis comparing the two groups: PS-I and I-PS. Figure 4a presents the distribution of relative learning gains within the two groups.

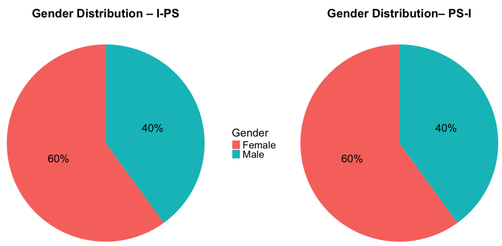

(a) Gender Distribution

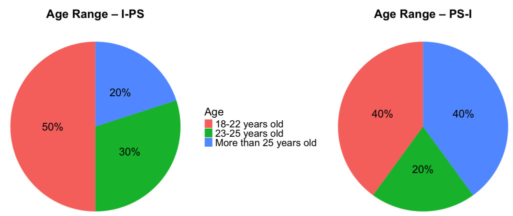

(b) Age Distribution

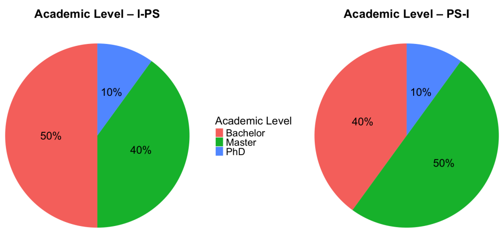

(c) Academic Level Distribution

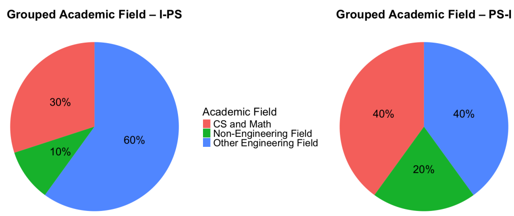

(d) Grouped Academic Field Distribution

Figure 3: Statistics of Participants

The PS-I group has a higher average relative learning gain (0.58) compared to the I-PS group (0.49); however, both groups display a broad range of learning gains. The I-PS group displays a wider spread in learning gain, and even includes a negative value. In contrast, the PS-I group's results are more tightly clustered toward the higher end of the scale, although there is still considerable variation in performance among individuals in this group. Figure 4b reveals that the overall distribution of relative learning gains is left-skewed, with a mode of 1, indicating that several students reached the maximum possible relative gain, 75% of which were from the PS-I group. This shows that some students performed exceptionally well, pulling the mean upward, consistent with the fact that for both groups, the mean is higher than the median.

Overall, the exploratory analysis revealed that the PS-I group tends to outperform the I-PS group in terms of average learning gains. However, the large within-group variation suggests that instructional method alone does not fully account for student performance. The observed skewness indicates that a small number of high performers may be disproportionately influencing the group averages. This highlights the need to explore additional contributing factors, as discussed in Section 4.2.

Group	**Mean**	**Median**	**Standard Deviation**
PS-I	0.58	0.53	0.33
I-PS	0.49	0.47	0.35

Table 1: Descriptive Statistics of Relative Learning Gain by Group

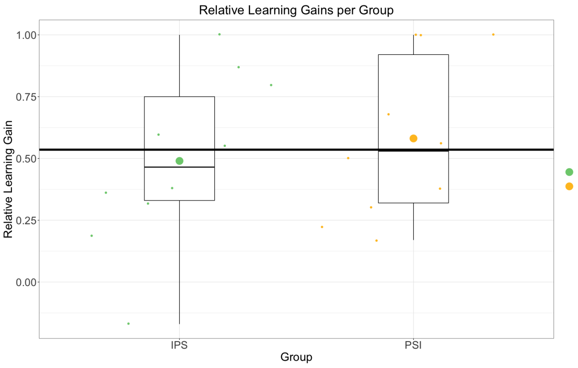

(a) Relative Learning Gains by Group

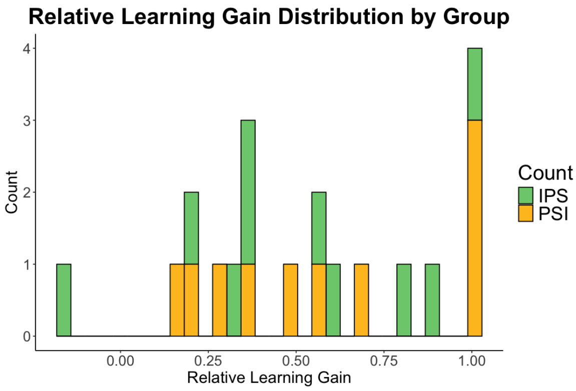

(b) Overall Learning Gain Distribution

Figure 4: Distributions of Learning Gains

6.2.2 Control Variables

Gender As observed in Figure 3a, gender is perfectly balanced between the two groups. Therefore, differences observed in the outcome are unlikely to be explained by this control variable.

6.2.3 Power Analysis

We conducted a power analysis to assess whether the sample size provides adequate statistical power to detect a significant effect of the independent variable (PS-I/I-PS) on the dependent variable (relative learning gain). We used `pwr.anova.test` function from package `pwr` in R with the following parameters: $\alpha = 0.05$ and cohen's f = 0.25 for a medium effect size. Using such parameters and with a sample size of 10 participants per group, we get a power of $\approx$ 18.51% (Type II error(β) $\approx$ 0.815), meaning that there is $\approx$ 18.51% probability of detecting an effect of instruction method on relative learning gain if it exists and a $\approx$ 81.49% chance of a false negative result. In order to detect a small to medium effect, with 95% power, we would require 210 participants.

6.2.4 One-way ANOVA

ANOVA assumptions A Shapiro-Wilk test on the model residuals yielded a $p - value$ of 0.43, and a Kolmogorov-Smirnov test had a p-value of 0.87, we therefore do not reject the null hypothesis of normality. A Bartlett test of homogeneity of variances yielded a $p - value$ of 0.85, so the same conclusion applies. We conclude that the residuals are approximately normally distributed and the variances are equal in both groups. The independence of observations was addressed by ensuring that each participant completed the activities individually and without collaboration. Therefore, we consider the assumption of independence to be reasonably satisfied, although we acknowledge the potential for shared background factors, as all participants were recruited through personal connections. We conducted a **one-way ANOVA** to test whether there is a significant difference in relative learning gain between the PS-I and I-PS groups. The results are summarized in Table 2.

Statistic	**Value**
Degrees of Freedom	18
F-statistic	0.36
p-value	0.56

Table 2: ANOVA Statistics for the Effect of Group on Relative Learning Gain

The model estimated the overall mean relative learning gain across groups to be approximately 0.54. The group effect coefficient indicated that the I-PS group scored about 0.045 lower than the overall mean, while the PS-I group scored about 0.045 higher. However, this difference was not statistically significant, with an F-statistic of 0.36 and a corresponding p-value of 0.56. Since p-value > 0.05, we fail to reject the null hypothesis of no difference in relative learning gain between the PS-I and I-PS groups.

6.2.5 Moderation Effects

For each moderation effect, we performed a two-way ANOVA to examine the impact of the effect on the relationship between the experiment group (PS-I or I-PS) and relative learning gain. Table 3 summarizes the moderation analyses conducted in this study. Figure 5 presents the plots of interaction effects between learning gains and the different moderators.

Table 3: Summary of Moderation Analysis (Two-way ANOVA)

Moderator	**Effect**	**F-statistic**	**df**	**p-value**	**Sample Size**	**Notes**
Confidence Level	Group	0.01	1,10	0.921	16	Low-frequency categories excluded
	Confidence	1.00	2,10	0.401		
	Interaction	0.49	2,10	0.624		
Academic Field	Group	0.051	1,14	0.824	20	**Significant main effect of field**
	Field	4.376	2,14	**0.033**		
	Interaction	0.119	2,14	0.889		
Academic Level	Group	0.08	1,14	0.775	20	No significant effects
	Level	1.01	2,14	0.388		
	Interaction	0.53	2,14	0.598		
Sports Level	Group	1.078	1,11	0.321	16	Low-frequency categories excluded
	Sports	0.447	1,11	0.517		
	Interaction	0.608	1,11	0.452		

Academic Field Effect The grouped academic field analysis revealed a significant main effect (F(2,14) = 4.376, p = 0.033), with CS and Math participants performing better in both PS-I and I-PS groups, as shown in Figure 5b.

Sample Size Reductions Several moderation analyses required sample size reductions due to low frequencies in certain categories. Confidence level analysis excluded "Not Confident at All" and "Very Confident" categories, reducing the sample from n=20 to n=16. This reduction may have contributed to lower statistical power in the two-way analysis. Similarly, sports level analysis excluded the ”0 times a week” category due to low sample size and unbalanced distribution, also reducing the sample to n=16.

Parents’ Education Level Distribution for parents’ education level was highly unbalanced (Primary school/High school/University: PS-I n=1/3/6, I-PS n=0/0/10). Due to this imbalance, we could not include this moderator in the analysis as meaningful comparison between education levels was not possible with two-way ANOVA.

Comparing Moderation Effect Magnitudes While none of the interaction effects achieved statistical significance, examining the relative magnitudes of F-values and p-values reveals meaningful patterns. Sports level demonstrated the strongest interaction potential ($F(1,11) = 0.608$, $p = 0.452$), followed by academic level ($F(2,14) = 0.53$, $p = 0.598$) and individual confidence level ($F(2,10) = 0.49$, $p = 0.624$). Academic field showed the weakest moderation effect, but was the only variable to show a significant main effect, suggesting that while academic background influences overall performance, it does not differentially affect the relative effectiveness of PS-I versus I-PS approaches.

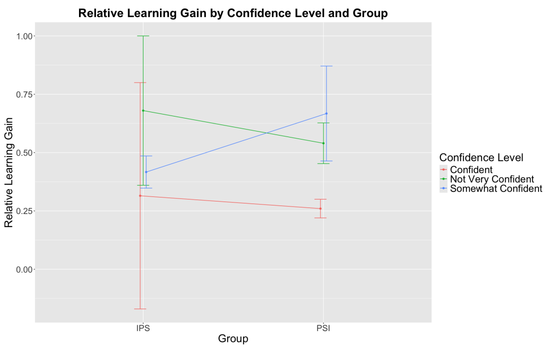

(a) Relative Learning Gains by Group and Confidence Level, Reduced Sample (n=16)

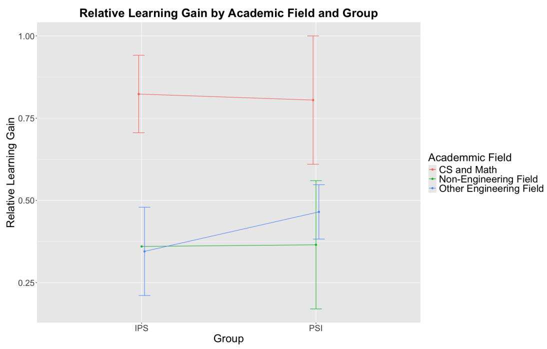

(b) Relative Learning Gain by Group and Grouped Academic Field

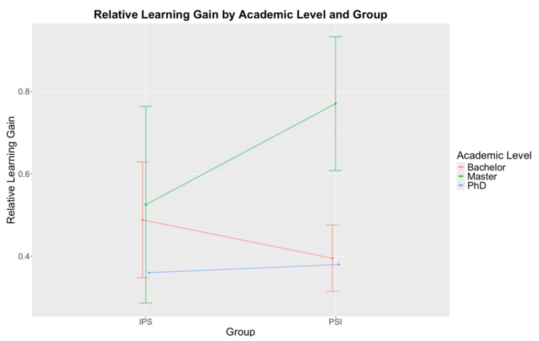

(c) Relative Learning Gain by Group and Academic Level

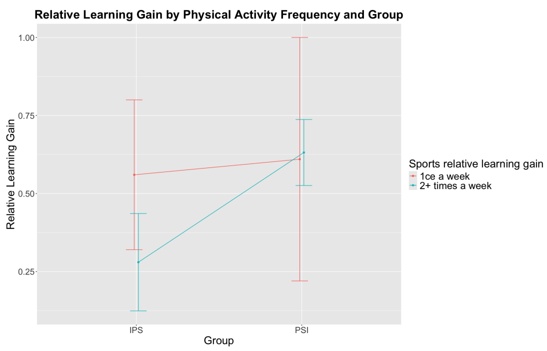

(d) Relative Learning Gain by Group and Physical Activity Frequency, (n=16)

Figure 5: Moderation Interaction Effects

6.2.6 Mediation effects

PS Algorithmic Solutions Analyzing the participants' PS activity solutions, we recorded two key solution types: algorithmic and non-algorithmic. We define non-algorithmic solutions as exhaustive solutions that blindly calculate the costs of all paths. Algorithmic ones, however, either follow a technique similar to Dijkstra or some logical assumptions, like pruning irrelevant paths. Figure 6a shows the impact of instruction condition (I-PS/PS-I) on solution type and the effect of the solution's type on learning gain. In I-PS, as participants perform the activity after watching the instruction video, there are more algorithmic solutions than PS-I ($a = -0.3$). We note a possible positive correlation between an algorithmic solution and relative learning gain, likely because an algorithmic solution reflects deeper engagement of the participant with the activity ($b = 0.16$). We record an average causal mediation effect (ACME) of -0.07. As PS-I can potentially have a higher impact on relative learning gain than I-PS (Figure 4a), such an ACME possibly suggests that PS-I decreases relative learning gain because of having less algorithmic solutions, and I-PS does not benefit from algorithmic solutions, as students just solve the activity using Dijkstra's algorithm taught in the video (i.e. algorithmic solutions in I-PS do not necessarily reflect deep engagement but rather lesson application). Our interpretation can be visually seen in Appendix A.3.1. We record an average direct effect (ADE) of 0.16, and none of our insights can be confirmed due to the high p-values.

PS Graphical Solutions We also categorize the participants' solutions according to whether they used visuals or not (i.e., representing train connections as a graph) and study the impact of a graphical solution on the relative learning gain. Similar to algorithmic solutions, graphical solutions are more common in I-PS than PS-I because participants see the graphical representation in the instruction video ($a = -0.3$). We note a small positive correlation between graph usage and relative learning ($b = 0.02$ and a small F-value of 0.02 with a high p-value of 0.88). We note a small ACME of -0.01, which may possibly show that using a graph helps to understand the problem better and engage more deeply with it in the case of PS-I. However, using a graph in I-PS is not as useful as it is just an application of the methodology taught in the instructional video. Finally, we note an ADE of 0.11 as in Figure 6b. For more visuals, see Appendix A.3.1.

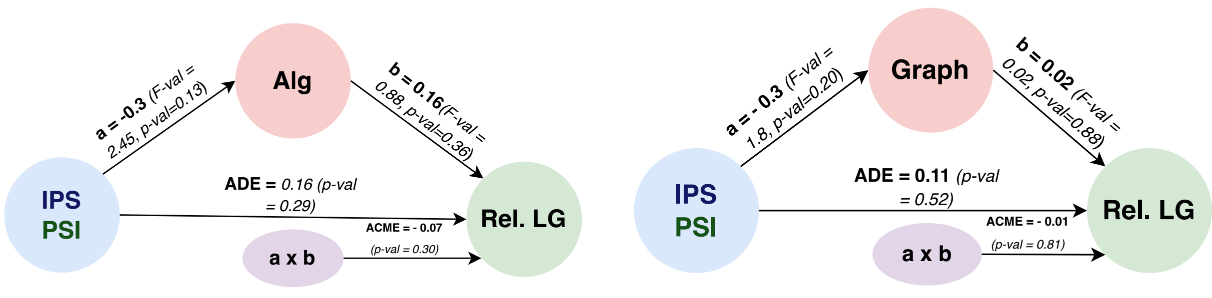

(a) Algorithmic versus Non-algorithmic

(b) Graphical versus Non-graphical

Figure 6: Mediation effect of PS activity solution type

Interest level In this analysis, we use the time spent in the pre-test and post-test as an indication of interest level (see § 4.2 for a detailed explanation of interest level measurement).

Using such a setup, we visualize (see Appendix A.3.1) a very insignificant impact of PS-I/I-PS on the interest level and of interest level on learning gain, so we don't show the mediation numbers.

6.2.7 Moderated Mediation effects

We wanted to investigate whether the possible mediators, algorithmic and graphical solution types, are moderated by academic field. In particular, we sought to answer this question: *"Are students from science fields more or less likely to produce algorithmic and graphical solutions?"*. Unfortunately, all the students we got for I-PS are from a scientific background, so we only provide visual insights from PS-I. First, we note that participants from a scientific background are more likely to produce graphical and algorithmic solutions than those from non-science fields (Figures 7a and 7b). Second, it is possible that the graphical solution type works better for science fields (*we can not confirm as there are only 2 non-science students in PSI*, Figure 8). There are no algorithmic solutions for non-science students, so we could not draw insights (Figure 7b).

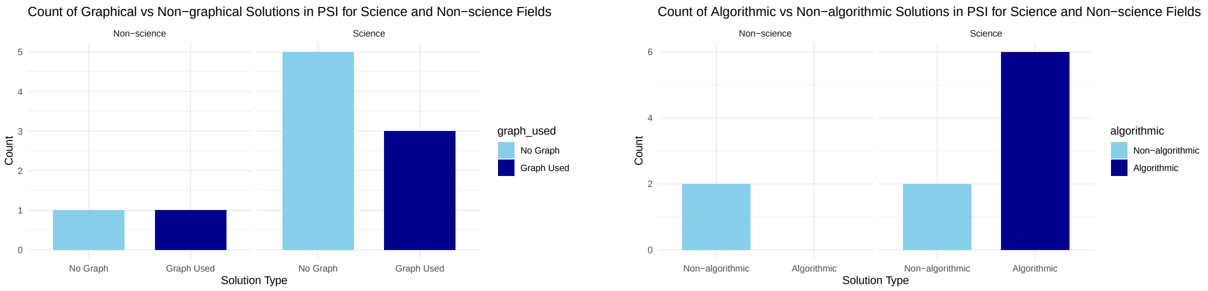

(a) Count of Graphical vs Non-graphical Solutions in PSI for Science and Non-science Fields

(b) Count of Algorithmic vs Non-algorithmic Solutions in PSI for Science and Non-science Fields

Figure 7: Distributions of Learning Gains

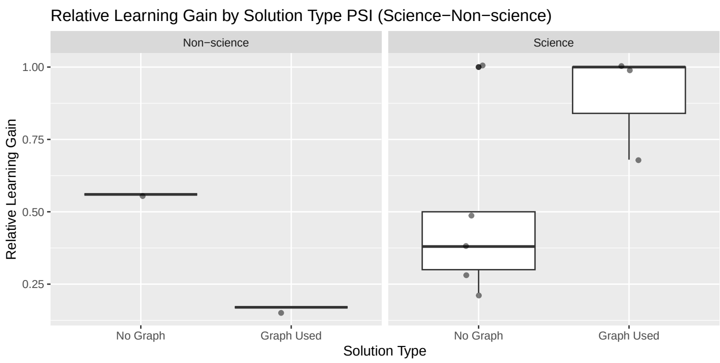

Figure 8: Relative Learning Gain by Solution Type PSI (Science vs Non-science).

7 Discussion and Conclusion

This study investigated whether Problem Solving followed by Instruction (PS-I) or Instruction followed by Problem Solving (I-PS) is more effective for learning Dijkstra's algorithm or not. While our analysis of 20 participants revealed no statistically significant difference between the two approaches (i.e., our test power is $\approx 18.51\%$), we found suggestive evidence that PS-I may offer advantages. In particular, our study provides the following insights. First, the PS-I group had a higher mean relative learning gain, with 75% of perfect scores coming from it. Second, students with backgrounds in Computer Science and Mathematics significantly outperformed those from other disciplines, regardless of instructional approach. This highlights the importance of considering academic background and prior knowledge when designing educational material. Third, while no interaction effects achieved statistical significance in our moderation analysis, the relative magnitudes of F and p-values suggest promising directions for future research. Sports level showed the strongest moderation potential, followed by academic level and confidence, showing that these factors may influence instructional approach effectiveness with adequate sample sizes. Finally, our mediation analysis suggests that algorithmic solution type may serve as a stronger mediator than graphical solution type in explaining the instructional approach effects, as the former has a much smaller p-value.

However, several limitations must be acknowledged when interpreting these findings. First, the small sample size of 20 participants restricts the statistical power of our analysis and reduces our ability to detect significant differences between the two instructional approaches, which may explain why the observed PS-I advantages did not reach statistical significance. An ad-hoc power analysis confirmed that our study was underpowered. Second, the recruitment of participants through personal connections may have introduced bias in the participant pool. Third, we initially wanted to recruit EPFL students (CMS, MAN or first years), who had not yet learnt about shortest path algorithms, as it would have been more insightful for this study. Because of recruiting challenges, however, we were forced to widen our search for participants to other fields and academic levels. Finally, the results reported may not be generalizable beyond participants who fit the profiles covered by our sample.

Despite these limitations, this study provides valuable preliminary evidence for the potential benefits of the PS-I approach and establishes a foundation for future research, with a larger, more representative sample to confirm the observed trends and explore the identified moderating and mediating factors.

References

[1] Tania Pinto-Escalona et al. "Sport Participation and Academic Performance in Young Elite Athletes". In: *International Journal of Environmental Research and Public Health* 19.23 (2022), p. 15651. DOI: 10.3390/ijerph192315651. URL: https://www.mdpi.com/1660-4601/19/23/15651.

[2] Jeff Johnson. "The Effect of Athletic Participation on Academic Achievement for Private High School Students". American English. PhD thesis. United States: Lynn University, Apr. 2023.

A Appendix

A.1 Lesson design

From	To	Travel Time
Kumasi	Wa	40 min
Tamale	Takoradi	4h
Takoradi	Accra	4h
Wa	Tamale	2h
Sunyani	Tamale	2h
Tamale	Accra	6h
Tamale	Cape Coast	2h
Wa	Sunyani	3h
Kumasi	Bolgatanga	3h 30min
Kumasi	Sunyani	4h
Sunyani	Accra	12h
Cape Coast	Accra	3h
Accra	Bolgatanga	20h
Bolgatanga	Bimbila	18h
Akosombo	Bimbila	22h
Kumasi	Akosombo	6h
Accra	Akosombo	16h

Figure 9: Train Table Given in the PS Activity

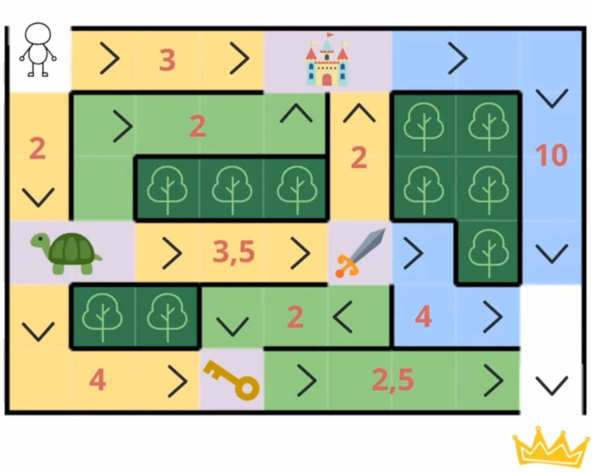

Figure 10: The Example Used in the I-PS Version of the Instructions with the Question: *What items does the character encounter to get to the crown in the shortest possible time?*

A.2 Implementation

A.2.1 Surveys

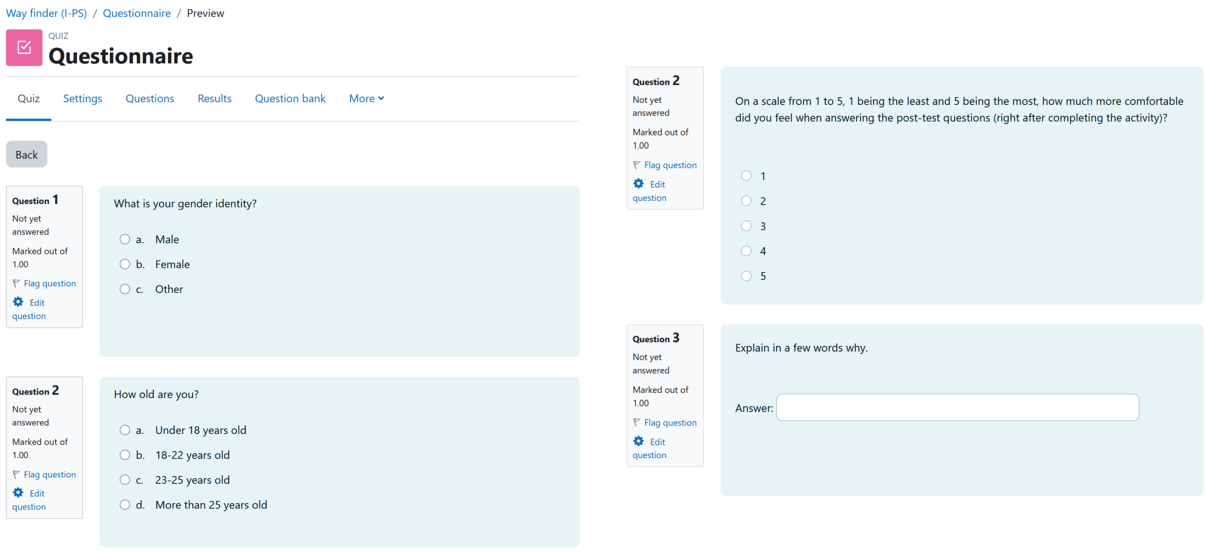

Figure 11: Moodle View of Some Questions of the Demographics (left) and Feedback (right) surveys.

A.2.2 Pre-test and Post-test Questions

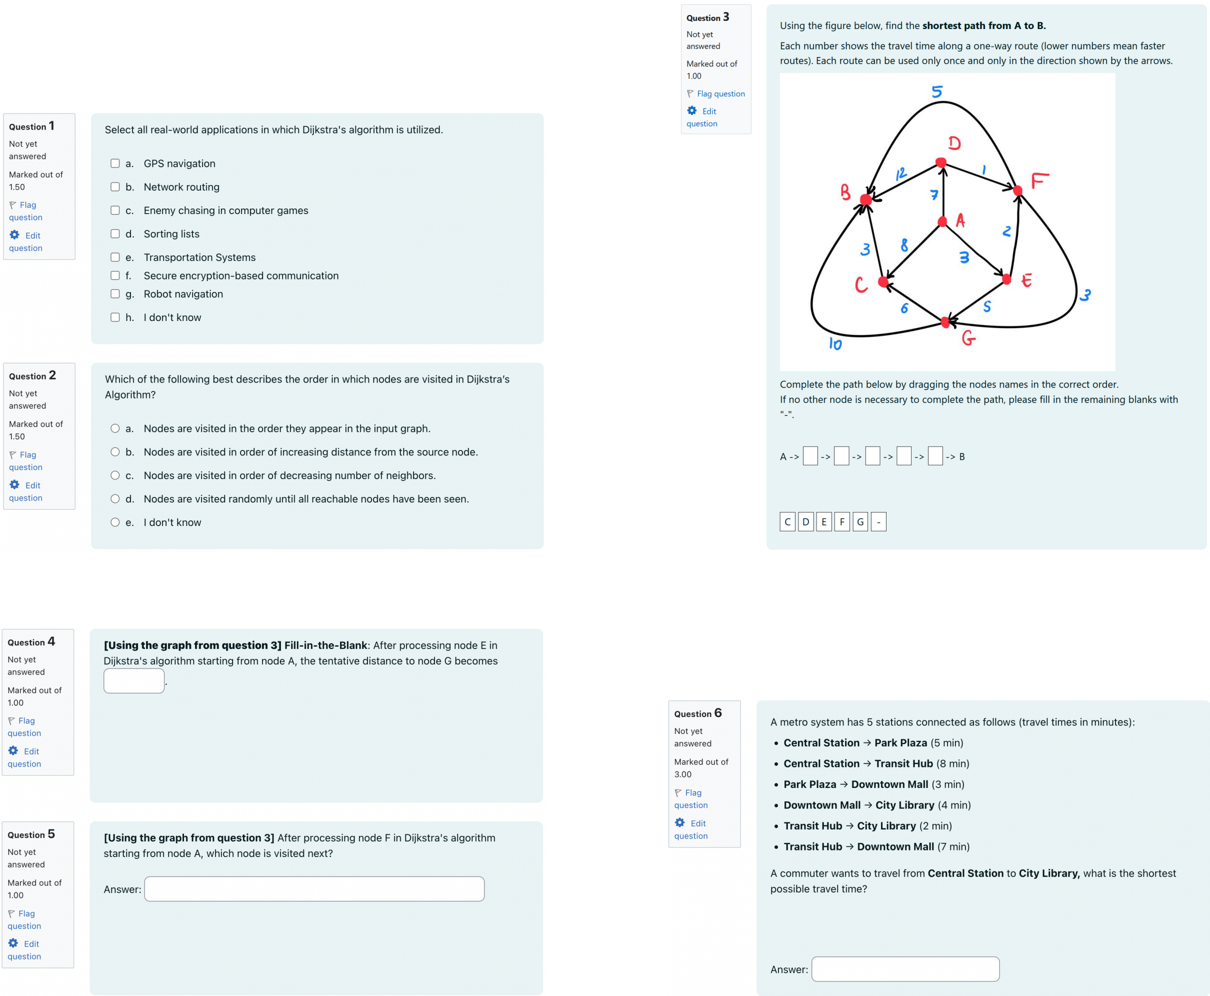

Figure 12: All Assessment Questions

A.2.3 PS Activity

Problem description and introduction The PS activity was introduced as follows:

Please complete this activity before proceeding to the next section. You will need a pen and a paper (and your phone to scan your work) or a tablet. Do not spend more than 20 minutes on this activity. It is completely fine if you don't find the correct solution. It's more important to show how you thought about the problem and how you tried to solve it.

Here is the full content of the activity as we presented it on Moodle:

```
Train networks connect cities with various travel times.  Finding the most
efficient route between distant cities often requires careful analysis
```

of multiple possible paths.

Train Timetable
Please note that the times in this example don't reflect the reality to avoid solving the problem based on how well you know the cities. The following table shows available train connections between different cities and their travel times:
See table in Figure 9 of Appendix A.1

Problem Statement
Francis needs to travel from Kumasi to Accra, departing at 2:00 PM. He must arrive in Accra before 10:00 PM on the same day.

Your Task
Determine which train connections Francis should take to reach Accra before 10:00 PM, starting from Kumasi at 2:00 PM.

Submission Instructions

1. Document your complete problem-solving process, showing:
   - How you approached the problem
   - Which possible routes you considered
   - How you calculated the total travel time or each option
   - Why you eliminated certain routes
2. You can sketch and write your reasoning on paper or tablet and upload your work below.
3. In the answer box, provide the final route in the following format: Kumasi → [city_1] → [city_2] → ... → [city_n] → Accra Where [city_1] through [city_n] represent the sequence of stops Francis must make.

Remember to verify that your solution allows Francis to arrive in Accra before 10:00 PM, considering all travel times between stops.

Screenshots

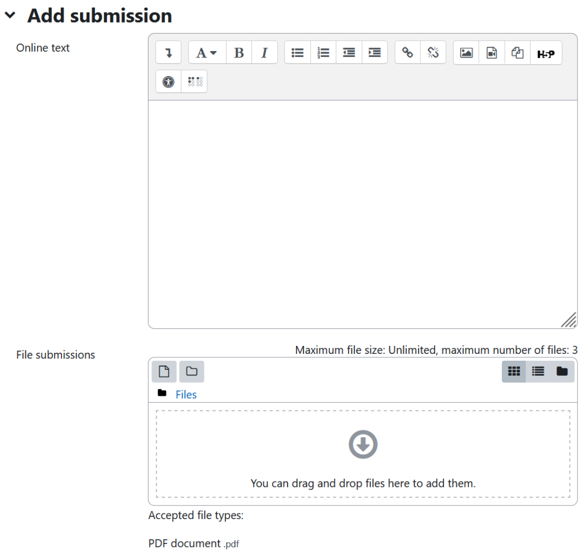

Figure 13: Submission Boxes at the End of PS activity

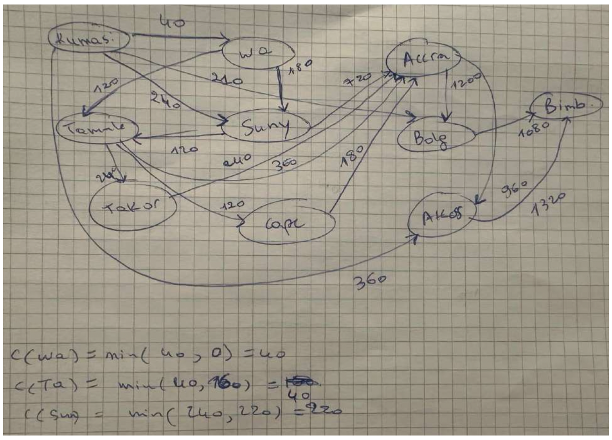

Figure 14: Part of a Submitted Solution in I-PS Problem Solving Activity

A.3 Data Analysis

A.3.1 Mediation Analysis

Solution Type (Algorithmic VS. Non-algorithmic

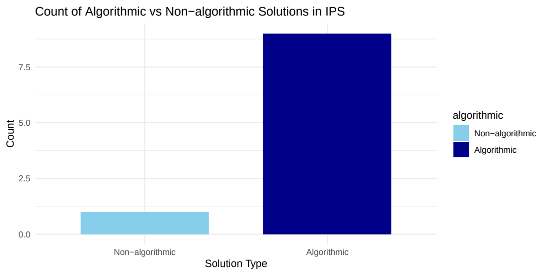

(a) Count of Algorithmic vs Non-algorithmic Solutions in I-PS

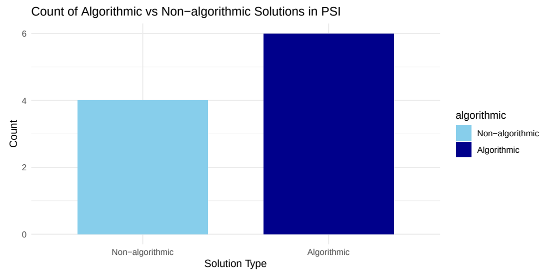

(b) Count of Algorithmic vs Non-algorithmic Solutions in PS-I

Figure 15: Comparison of Algorithmic and Non-algorithmic Solutions in I-PS and PSI

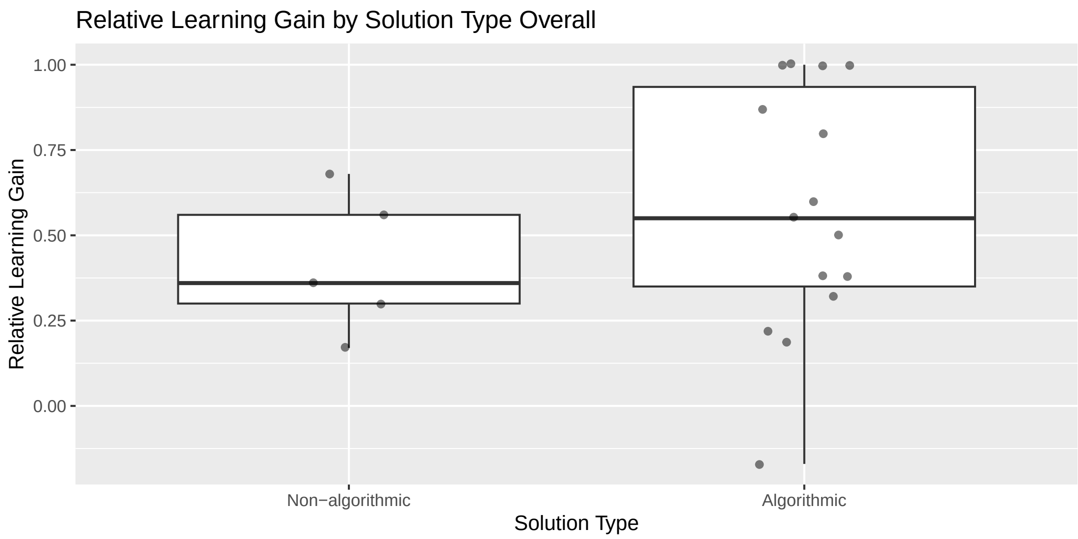

Figure 16: Overall Relative Learning Gain by Solution Type (Algorithmic vs Non-algorithmic)

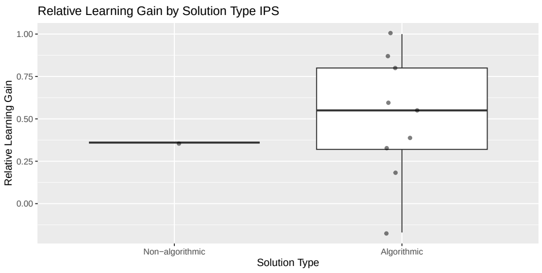

(a) I-PS Relative Learning Gain by Solution Type (Algorithmic vs Non-algorithmic)

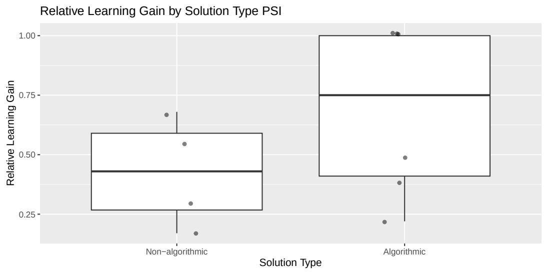

(b) PS-I Relative Learning Gain by Solution Type (Algorithmic vs Non-algorithmic)

Figure 17: Comparison of Relative Learning Gain by Solution Type for I-PS and PS-I

Solution Type (Graphical VS. Non-graphical

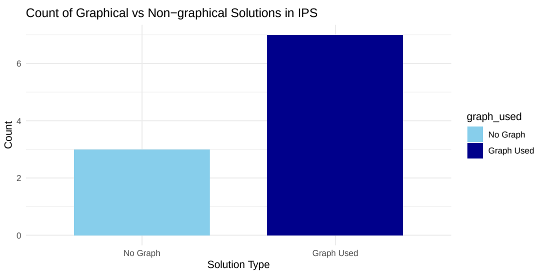

(a) Count of Graphical vs Non-graphical Solutions in I-PS

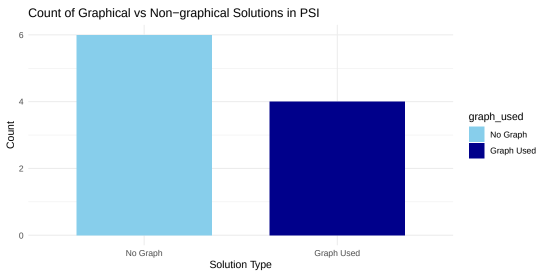

(b) Count of Graphical vs Non-graphical Solutions in PS-I

Figure 18: Comparison of Graphical and Non-graphical Solutions in I-PS and PS-I

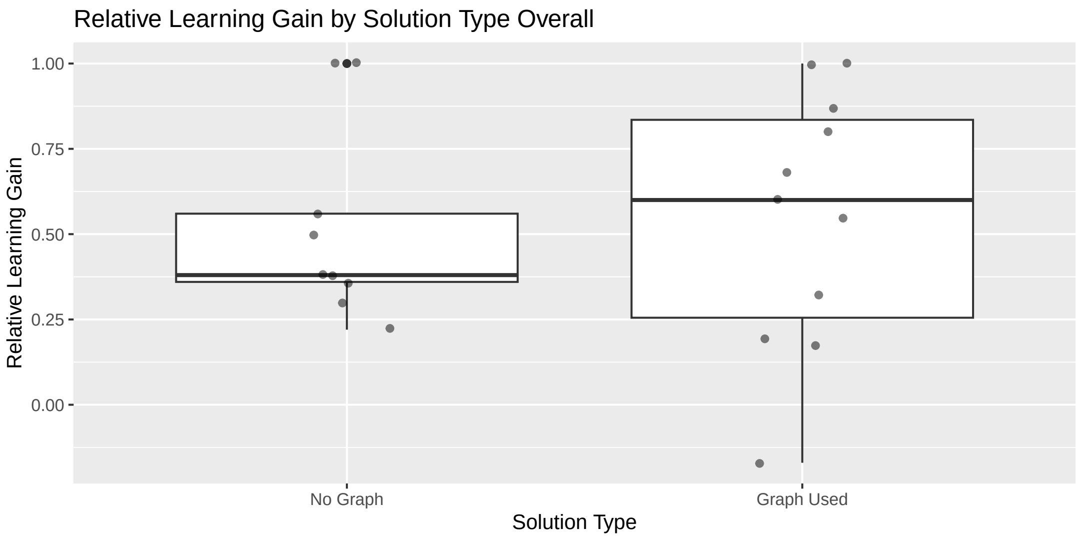

Figure 19: Overall Relative Learning Gain by Solution Type (Graphical vs Non-graphical)

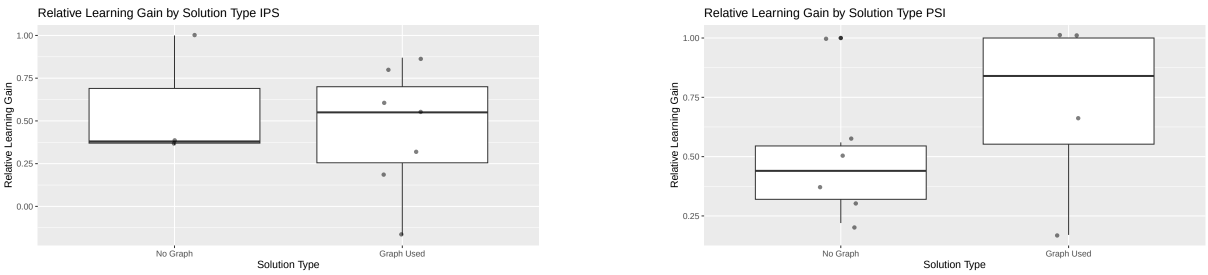

(a) I-PS Relative Learning Gain by Solution Type (Graphical vs Non-graphical)

(b) PS-I Relative Learning Gain by Solution Type (Graphical vs Non-graphical)

Figure 20: Comparison of Relative Learning Gain by Solution Type (Graphical vs Non-graphical) for I-PS and PS-I

Interest Level

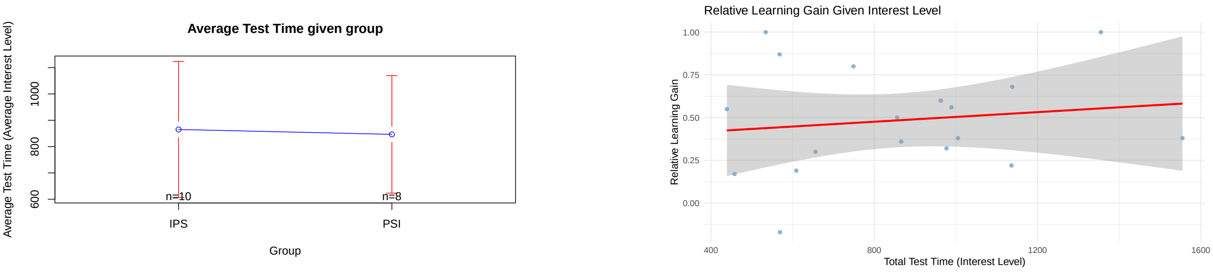

(a) Average Test Time given group

(b) Relative Learning Gain Given Interest Level

Figure 21: Comparison of Average Test Time and Relative Learning Gain