

Iguanodon: A Code-Breaking Game for Improving Visualization Literacy

Patrick Adelberger , Oleg Lesota , Klaus Eckelt , Markus Schedl , and Marc Streit 

Abstract—In today’s data-rich environment, visualization literacy—the ability to understand and communicate information through visualizations—is increasingly important. However, constructing effective visualizations can be challenging due to the vast design choices involved. Off-the-shelf systems and libraries produce visualizations with carefully selected defaults that users may not be aware of, making it hard to increase their visualization literacy with those systems. In addition, traditional ways of improving visualization literacy, such as textbooks and tutorials, can be burdensome as they require sifting through a plethora of resources. To address this challenge, we designed Iguanodon, an easy-to-use game-like application that complements the traditional methods of improving visualization literacy. In our game application, users interactively choose whether to apply design choices, which we assign to sub-tasks that must be optimized to create an effective visualization. The application offers multiple game variations to help users learn how different design choices should be applied for constructing effective visualizations. Furthermore, our approach easily adapts to different visualization design guidelines. We describe the application’s design and present the results of a user study with 37 participants. Our findings indicate that our game-based approach supports users in improving their visualization literacy.

Index Terms—Gamification, serious games, visualization, education, learning

1 INTRODUCTION

THE use of visualizations has become more prevalent in recent years. It can now be found frequently in various forms, for example, in public media, lifestyle applications, and professional software. Hence, proficiency in reading and creating visualizations is a crucial skill that is essential not only in specific fields such as science and business but also in everyday life. The ability to confidently create and interpret visual representations of data is known as visualization literacy. Solen [1] defines visualization literacy as “the ability to critically interpret and construct engaging visualizations.”

Constructing effective visualizations is a complex task that demands a considerable amount of knowledge and expertise. Novices may not have the necessary skills to create high-quality visualizations, requiring them to invest time in acquiring relevant knowledge. To enhance their visualization literacy, they may need to refer to a wide range of learning resources, such as textbooks and tutorials. By developing a high level of visualization literacy and gaining a better understanding of visualization design principles, novices can learn how to effectively create compelling and informative visualizations, while also improving their ability to read and interpret visualizations created by others.

Off-the-shelf visualization tools like Tableau [2] and PowerBI [3] provide meaningful defaults for the various design choices to support users in creating effective visualizations. Although tools that offer ready-to-use visualizations or recommend visual representations based on the loaded dataset can reduce the effort required, they do not reveal the

design choices that underlie the visualizations. This implicit knowledge diminishes the learning effect of such tools.

A possible way to make learning more engaging is the use of serious games—computer games that were designed to help users learn various aspects, such as mathematics, health, etc. Connolly et al. [4] offer a comprehensive overview of serious games and their evaluations in game-based learning. Their findings demonstrate the efficacy of game-based learning, supported by empirical evidence.

Against this background, our work makes two contributions. First, a serious game application named Iguanodon that breaks down the challenging task of creating effective visualizations into smaller, more manageable sub-tasks. These tasks must be satisfied by applying or not applying design choices to a visualization. This concept makes it easy to adapt our serious game to different visualization design guidelines. Second, a user study evaluating the effectiveness of our approach in improving users’ visualization creation skills. As an auxiliary contribution, we present a visualization literacy test that focuses on the construction aspect for assessing the visualization literacy of study participants.

2 RELATED WORK

In this section, we first discuss the importance of visualization literacy and ways to assess it. We then give a short overview of recommendation systems that support users in creating effective visualizations without necessarily being visually literate. Finally, we elaborate on different tools and game applications and their approaches to improving visualization literacy.

2.1 Visualization Literacy

The work by Solen [1] reviews existing literature on visualization literacy and provides a list of the various definitions

• All authors are with the Johannes Kepler University Linz. E-Mail: {patrick.adelberger, oleg.lesota, klaus.eckelt, markus.schedl, marc.streit}@jku.at.

Manuscript received MONTH DAY, 2023; revised MONTH DAY, 2023.

available. Lee et al. [5] defined visualization literacy as “the ability and skill to read and interpret visually represented data and extract information from data visualizations”. Börner et al. [6] defined it as “the ability to make meaning from and interpret patterns, trends, and correlations in visual representations of data”. While these definitions focus on the reading and interpretation aspect, Boy et al. [7] additionally included the construction of visualizations in their definition. In our work, we rely on the visualization literacy definition by Solen [1] and define it as the ability to critically interpret and construct visualizations.

Solen [1] surveyed the different participant pools of previous studies on visualization literacy, and found that user studies are mostly limited to the Western, educated, prosperous, and democratic population. Furthermore, they proposed an extended definition of visualization literacy that includes construction, interpretation, believability, and engagement. Believability stands for topics such as trust and bias. In their survey paper and the follow-up work, Firat et al. [8], [9] classified the methods used to evaluate visualization literacy. They presented an overview of the evaluation types, the target audience, and the number of participants. They found that within-subject design studies are more prevalent and usually involve participants with high school or higher education levels. In addition, they reported bar charts, scatterplots, and line charts to be the most commonly evaluated visualization types in the context of visualization literacy research. However, from the works mentioned above, it can be concluded that no standard is established for assessing visualization literacy, neither for reading nor for creating visualizations.

Börner et al. [6] assessed the visualization literacy of 273 science museum visitors with 20 visualizations from four visualization types (chart, graph, map, network layouts). They found that a large share of their participants could not name or interpret the different visualization types commonly used in newspapers, textbooks, or on the internet. The visualization literacy test developed by Boy et al. [7] is based on item response theory to consider the test items’ difficulty, covering six different tasks: maximum, minimum, variation, intersection, average, and comparison. They created four tests for three visualization types (line graph, bar chart, and scatterplot) to quickly assess the participants’ visualization literacy. A more extensive and detailed test to measure visualization literacy is the visualization literacy assessment test (VLAT) proposed by Lee et al. [5]. It consists of 53 multiple-choice questions for twelve visualization types, including line charts, bar charts, and maps. The questions cover eight tasks, such as retrieving value, finding extremum, and making comparisons.

Börner et al. [10] presented a data visualization framework that provides teaching exercises and assessments for constructing and interpreting data visualizations. This framework was developed over the last decades and includes a topology of core concepts and a process model. According to the authors, the combination of the topology and the process steps can be used to design teaching materials, exercises, and visualization literacy assessment tests.

For the purpose of our work, we developed our own assessment test that focuses on evaluating the construction component of visualization literacy. We designed our test

in a manner that allows us to evaluate the efficiency of our game application, as justified in Section 4.1.

2.2 Visualization Creation and Recommendation

Creating effective visualizations is a challenging task for individuals with limited experience or understanding of data visualization. However, visualization recommendation systems can assist users in creating visualizations by suggesting what data subset to use and how it should be visualized. There are three main approaches these tools use to create visualizations: rule-based, machine learning (ML)-based, and approaches that do not fit these categories.

Systems that follow a rule-based approach rely on predefined rules to determine the most appropriate visualization for a particular data type or context. An early example is the Show Me [11] functionality integrated into Tableau [12]. Users choose one or multiple attributes of interest and Show Me suggests suitable visualizations. After selecting the visualization, users can refine them in Tableau. Voyager [13] and Voyager 2 [14] not only allow the creation of visualizations but also recommend further visualizations based on statistical and perceptual measures. VizAct [15] supports users with step-by-step guidance through the creation of visualizations and tracks their actions during the process. Users can later review or revisit these tracked actions.

Systems that adopt an ML-based approach use machine learning algorithms to analyze the characteristics of a given dataset and automatically recommend suitable visualizations. VisML [16], for example, uses a neural network trained on a large number of datasets and their corresponding visualizations to recommend visualizations of a given dataset. The Data2Vis [17] system creates visualization recommendations based on a training set containing Vega-Lite visualizations [18].

Some systems, however, do not fit into the categories mentioned above. DeepEye [19] and Draco [20], for example, use a hybrid approach combining rule-based and ML-based approaches. DeepEye enhances its ranking of recommended visualizations by integrating additional feedback from experts, who create rules based on their domain knowledge to rank visualizations by their quality. Draco is a constraint-based system that represents design knowledge as constraints with weights. These weights can be learned by a ranking model trained on ranked pairs of visualizations. The KG4Vis [21] system generates a knowledge graph by extracting triplets—representing a data column’s features and visualization design choices—from pairs of datasets and visualizations. Using this approach, the system can automatically recommend visualizations and derive rules based on these triplets enabling users to comprehend the recommendation results.

These visualization tools showcase the complexity involved in creating compelling visualizations and the significant effort that goes into developing them. Additionally, many of these systems use meaningful default design choices that users may not be aware of. Therefore, to create effective visualizations, it is essential for individuals to improve their visualization literacy and gain a deeper understanding of the design choices made by these systems. This knowledge will help users make informed decisions

when creating their own visualizations, even without these tools.

In cases where users create visualizations manually by writing code, tools such as VizLinter [22] and VisuaLint [23] highlight errors in the visualization that need to be addressed to improve the visualization. VizLinter additionally offers information about the mistakes made and ways of fixing them. Similarly, VisuaLint identifies errors directly in the visualization by highlighting the problematic elements in red together with further information about the problems.

2.3 Visualization Literacy Games and Tools

Different ways exist to improve visualization literacy. The traditional approach is to learn through books, courses, tutorials, and guides, which provide a solid foundation for understanding the principles of data visualization. An alternative approach is to utilize games and tools specifically designed to increase visualization literacy.

Alper et al. [24] and Bishop et al. [25] studied how children between the ages of five and eleven construct visualizations. They analyzed how children use drag-and-drop interactions to build visualizations with different abstraction levels, from free-form pictographs to bar charts. Diagram Safari [26] is a game where children learn to create, read, and interpret bar charts. Players have to create bar charts in such a way that an armadillo can roll from the start position to the goal. Huynh et al. [27] designed a story-based game for children to improve visualization literacy for pie charts and histograms. Players are introduced to their main characters and have to meet other characters to help them answer visualization questions.

VizItCards [28] is a toolkit that helps students improve their visualization literacy. The core components of this toolkit are cards, which can be grouped into different types. These types include cards showing what a domain user would do, actions to apply to a dataset, or different visualization types. Variants of the game cover a diverse set of domains, datasets, and goals. The variants can be split into two types: deconstruction and construction. Deconstruction is a top-down approach where one starts with a visualization and has to identify all its components. In contrast, construction is a bottom-up approach where one builds a visualization by combining visual elements. Amabili et al. [29] also define the two models—deconstruction and construction—in their taxonomy for creating educational games and showcase each model with a card game. They created two sets of cards: *sliding cards* that show a visualization on the top layer and its characteristics on the bottom layer and *legend cards* that depict various visualization characteristics. Based on the deconstruction model, they designed the Guess Viz? game in which each player selects a sliding card with a visual representation and has to ask questions to guess the visualization of the other. In the construction-based From A to viZ game, players pick legend cards that represent visualization requirements, for example, a visualization with nominal data for non-experts. The player who comes up with the most visualizations satisfying the requirements wins.

In our work, we designed a game application to enhance visualization literacy by emphasizing the construction process. This enables users to learn how to apply established

design choices to visualizations. We defined several game variations to illustrate the adaptability of our application for different visualization design guidelines.

3 IGUANODON DESIGN

Constructing visualizations is a complex task that requires careful consideration of many design choices. For our game application, we split this high-level goal of creating an effective visualization into multiple *objectives* that users need to address. In Iguanodon users can take *actions* towards fulfilling the objectives. Actions reflect design choices taken from literature. We designed these actions as binary operations, meaning an action can be either applied or not applied to a visualization. The visualizations depict the combinations of the differently applied actions. Additionally, a change in an action's state is directly applied to the visualization, resulting in an interactive learning experience. Each action maps to one objective, but multiple actions can address an objective.





To demonstrate the capabilities of Iguanodon, we created three game variations, each using objectives and actions based on published books, papers, or other resources. Our design for these game variations focuses on both general actions that may be part of the default design choices of some creation and recommendation systems discussed in Section 2.2, as well as specific actions tailored to a specific visualization type. To assess the effectiveness of Iguanodon, we restricted the visualization to a scatterplot and actions associated with this type of visualization for all game variations. The goal of our work was not to provide an exhaustive list of objectives and actions for various visualization types but to show that our approach helps users improve their proficiency in creating visualizations, which can then be applied to other visualization types as well.












































To determine the game variations and their corresponding objectives and actions, we started by referring to Rees and Laramee's survey on information visualization books [30]. The survey provided us with a starting point to identify relevant published resources that could be used for our game variations. After careful consideration, we have decided to base two out of the three game variations on individual books. The first game variation is based on the objectives and actions mentioned in *The Visual Display of Quantitative Information* by Tufte [31]. In the second game variation, we adopted the guidelines presented in *Show Me the Numbers by Few* [32]. In contrast, the third game variation incorporates all objectives and actions extracted from both books and adds additional actions identified in the literature to the existing objectives. We call the three game variations the *Tufte Game*, *Few Game*, and *Mixed Game*, respectively.







3.1 Objectives and Actions

In this section, we introduce the extracted objectives and their actions based on published work. Additionally, we add further resources for the objectives and actions already extracted from Tufte's [31] and Few's [32] books. The overview in Table 1 shows the game variations with their objectives and actions in their initial state.




TABLE 1




Overview of all game variations in their initial state with their actions and objectives.  and  represents that the action with its corresponding objective is included in the game variation.  indicates and applied action and  a not applied action.

Objective	Action	Tufte Game	Few Game	Mixed Game
 O1 Reduce overplotting	 A1.1 Decreased mark size			
	 A1.2 Change mark type to ring	-		
	 A1.3 Decrease mark opacity	-	-	
	 A1.4 Sample data randomly (25%)	-	-	
	 A1.5 Aggregate data points (mean)	-	-	
 O2 Reduce chartjunk	 A2.1 Add background			
	 A2.2 Lighten grid lines			
 O3 Improve readability	 A3.1 Use horizontal text for y-axis title			
	 A3.2 Write abbreviations out		-	
 O4 Utilize color encoding properly	 A4.1 Add legend	-		
	 A4.2 Apply nominal color schema	-		
	 A4.3 Add border around legend	-		
	 A4.4 Add legend title	-	-	

 **O1** *Reduce overplotting*. A common problem in visualization in general, and with scatterplots in particular, is overplotting, where too many data points are plotted on top of each other. In such cases, it can be difficult to judge the distribution of points correctly. This is the case for all game variations. This objective includes all actions that help mitigate the problem of overplotting in scatterplots. Tufte states that the data measures must shrink with a larger amount of data items, for example, smaller dots for scatterplots [31, p. 168]. Similar suggestions for reducing the mark size to reduce overplotting can be found in the resources by Few [32, p. 209], Ellis and Dix [33], and Camões [34], Holtz [35]. We defined the action  **A1.1** *Decrease mark size* that reduces the overlap of the data items in the scatterplot when applied. Few [32, p. 210] and Camões [34] proposes another option to reduce overplotting. They suggest not filling out the data point mark items and only using their border, resulting in a ring. We adapted this suggestion in the  **A1.2** *Change mark type to ring* action, where a ring replaces the point data mark item. Ellis and Dix [33] discuss further options to mitigate overplotting. One such option is reducing the opacity of the marks to see the density of items in a region. This action is also suggested by Camões [34], Holtz [35], and Kirk [36]. We reflected this suggestion as action  **A1.3** *Decrease mark opacity* that makes the data points transparent, using an opacity value of 50%. Another effective method, presented by Ellis and Dix [33], to reduce overplotting is to sample the data by selecting a subset of data points, thus decreasing overlap. We added this suggestion as an action and implemented  **A1.4** *Sample data randomly (25%)*, which randomly selects 25% of the data items. The method of sampling the data to reduce overlap is also recommended by Keim [37] and Holtz [35]. Keim [37] suggested two methods to mitigate overplotting. First, sampling the data, as mentioned, and second, aggregating the data. We have added the action  **A1.5** *Aggregate data points (mean)*, which calculates the mean value of the data points. In cases where color coding is used, we add a point for each attribute value used for the color encoding. This point represents the mean of the items

grouped based on this attribute value.

 **O2** *Reduce chartjunk*. According to Tufte, visualizations should minimize the use of ink for non-data-related elements. He defines elements not related to the data as *chartjunk* [31, p. 121]. We have taken Tufte’s definition of chartjunk and mapped two actions targeting this objective based on his recommendations. The first action  **A2.1** *Add background* adds a gray background to the visualization. A colored background could either distract from the visualization or reduce contrast, making it harder to read the data. According to Few, a visualization should not have a colored background to ensure that the data is visually prominent [32, p. 251]. Franconeri et al. [38] also suggest this in their work.  **A2.2** *Lighten grid lines* is based on Tufte’s recommendation to lighten the grid lines in a chart. Similarly, Few points out that the grid lines are support components and should therefore be kept light [32, p. 241]. Consequently, we show the grid lines in light gray to make them less distracting and allow the data to stand out more clearly. This action is also recommended by Franconeri et al. [38], and the *Data Visualization Standards* [39].

 **O3** *Improve readability*. To improve the readability of visualizations, we have adopted Tufte’s *Friendly Data Graphic* [31, p. 183] example. This objective aims to make it easier for readers to understand the information presented in the visualization. We specified the action  **A3.1** *Use horizontal text for y-axis title*, which changes the axis title from vertical to horizontal, i.e. the reading direction in the English language, and places it at the top of the axis to save space. This is also reflected throughout Few’s textbook, where he uses horizontal y-axis titles placed at the top for all visualizations. Additionally, we mapped the  **A3.2** *Write abbreviations out* action to this objective. This action ensures meaningful labels without abbreviated text that could confuse users.


 **O4** *Utilize color encoding properly*. We offer color encoding in our visualization for the Few and the Mixed Game. We created this objective to consolidate all the actions related to color encoding. According to Few, if a color encoding is used, it should be properly labeled [32, p. 244].



Fig. 1. The game interface of Iguanodon. Two visualizations **A** show the previous and current attempts of the players to *break the code*. Players have up to 10 attempts to improve the visualizations by applying various actions **B**. The heatmap **C** summarizes how well the design goals have already been met. The toolbar at the top allows players to select a game variation (Tufte, Few, Mixed) and provides access to help and onboarding information.

The conventional way to label color encoding is with a legend, which can be added using the **A4.1** *Add legend* action. This is also recommended by the *Data Visualization Standards* [39]. Few [32, p. 344], Camões [34], Kirk [36], and Franconeri et al. [38] provide examples of which type of color schema should be used for categorical data. We applied this guideline using the **A4.2** *Apply nominal color schema* action, where different hues are used to color the various categorical data values. Furthermore, Few cautions against using a border around the legend, as it can draw unnecessary attention to it [32, p. 246]. To represent this guideline, we included an action that adds a border around the legend: **A4.3** *Add border around legend*. The *Data Visualization Standards* [39] recommend using a legend title to add additional context for the color-encoded data. We integrated this recommendation as action **A4.4** *Add legend title* and mapped it to this objective.



3.2 Interface and Gameplay























































































Iguanodon's game interface is inspired by the popular code-breaking board game Mastermind [40]. In Mastermind, a solution is hidden, and by guessing how to apply different combinations of *code pegs*, the users can solve the code and win the game. The users have multiple attempts and get feedback on how close they are to the solution after each attempt. Guesses and feedback alternate until the users run out of guess attempts or win the game. In contrast to Mastermind, Iguanodon represents the solution, code pegs, and feedback as an effective visualization, actions, and objectives' state, respectively.

Iguanodon's interface consists of two parts stacked on top of each other: users work with the actions and a visualization to solve a game. Therefore the upper part of the interface displays the visualization based on the current configurations of the users' selected actions' state. This allows users to visually judge the effectiveness of the current visualization, as shown in Figure 1 **A**. Additionally, the visualization of the previous attempt is displayed on the left of the current visualization as a reference for the next attempt. The lower part of the interface displays the modified version of the Mastermind game. All available actions for the game variation are listed as toggle buttons, and users can apply () or not apply () them to modify the visualization for the current attempt, as shown in Figure 1 **B**. If users believe the current visualization is a potential solution to the game, they can confirm the current attempt, and the state of the objectives will provide feedback, as shown in Figure 1 **C**. The objectives can be either *not fulfilled* (Red), *partially fulfilled* (Yellow), or *fulfilled* (Green), depending on the correctness of the current visualization. Table 2 provides an overview of the different game variations' solutions with the fulfillment criteria of the objectives. The actions of the **O1** *Reduce overplotting* objective interact with each other, which is why there are several possible solutions for the Few Game and the Mixed Game. The users have ten attempts per game to solve it. However, if users cannot create a satisfactory visualization within ten attempts, a solution visualization with all the states of its actions will be revealed. Thus, users can compare their attempts with the solution.

We included a scoring system similar to Mastermind to make the game application more engaging and enable users

TABLE 2

Overview of the different solutions for each game variation and their objectives: **O1** Reduce overplotting, **O2** Reduce chartjunk, **O3** Improve readability, **O4** Utilize color encoding properly. An action marked with  indicates that the action is applied to the visualization and  represents an action not applied.

Objective	Action	Solutions for the different game variations									
		Tufte			Few			Mixed			
		#1	#1	#2	#1	#2	#3	#4	#5	#6	#7
O1	A1.1 Decreased mark size										
	A1.2 Change mark type to ring	-									
	A1.3 Decrease mark opacity	-	-	-							
	A1.4 Sample data randomly (25%)	-	-	-							
	A1.5 Aggregate data points (mean)	-	-	-							
O2	A2.1 Add background										
	A2.2 Lighten grid lines										
O3	A3.1 Use horizontal text for y-axis title										
	A3.2 Write abbreviations out		-	-							
O4	A4.1 Add legend	-									
	A4.2 Apply nominal color schema	-									
	A4.3 Add border around legend	-									
	A4.4 Add legend title	-	-	-							

to track their results. The score is based on the number of attempts required to solve the game. Therefore, the fewer attempts are needed, the higher the score. To motivate users to achieve better results, we introduced three reward badges—gold (solved in one attempt), silver (two attempts), and bronze (three attempts).

We added a menu bar at the top of the interface. With the *Help* button, users are presented with onboarding information about the application’s goal, application elements, how to play it, and the dataset. All game variations use the cars dataset by the American Statistical Association (ASA) [41] and apply the same data mapping on the axes of the scatterplot. The x-axis shows the horsepower of the cars, and the y-axis depicts the miles per gallon. Depending on the game variation, we color-code the cars by country of origin. For each visualization, we provide additional information about the dataset and its size.

Furthermore, users can restart the current game anytime by clicking the *Retry* button or switch to a different game variation using the drop-down in the menu bar. This drop-down also includes the current high scores for the different game variations.

Iguanodon’s approach makes it easy for users to learn how to create effective visualizations, even if they are novices in the field. The linkage between the action changes and the visualization allows users to see an action’s impact directly.

3.3 Implementation

The Iguanodon application is deployed at <https://jku-vds-lab.at/iguanodon>. The open-source code is available on GitHub: <https://github.com/jku-vds-lab/iguanodon>. The web application is written in Typescript, uses Vega-Lite [18] for all visualizations, and arquerio [42] for the dataset handling.

4 EVALUATION

We conducted a within-subject user study with a pre- and post-game test to evaluate whether our game application can enhance users’ proficiency in creating visualizations.

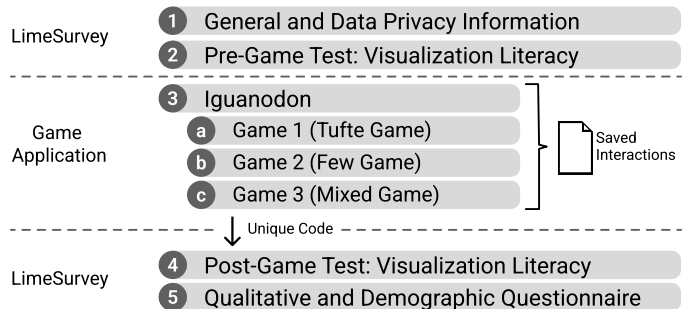


Fig. 2. The study structure with its five steps. The game application generates a code after all necessary game variations are played. The interactions in the game application are tracked and saved on a server. The saved interactions are matched to the LimeSurvey study data via the code.

We conducted the user study using the open-source survey tool LimeSurvey [43]. The study consisted of five parts, as illustrated in Figure 2. At the beginning of the study, the participants were given an introduction that explained the purpose of the user study and included information about data protection. After the introduction, we required participants to take a visualization literacy test (see Section 4.1 for more detail) to assess their existing knowledge. We refer to this first visualization literacy test as the pre-game test. After completing the pre-game test, participants were directed to the game application, where they had to solve the three game variations described in Section 3. To ensure that all participants were familiar with all objectives and actions, we asked them to play each game variation at least once

TABLE 3

Overview of the hypotheses tested in our user study. ✓ indicates an accepted hypothesis and - indicates a rejected hypothesis. The four different objectives are: **O1** Reduce overplotting, **O2** Reduce chartjunk, **O3** Improve readability, **O4** Utilize color encoding properly. Experts are participants with a (self-reported) lot of experience and non-experts have some or no (self-reported) experience.

#	Hypothesis	Accept
H0 ₀	There is no statistical difference between the overall pre-game and post-game test scores.	-
H0 _A	The overall post-game test score is statistically significantly higher than the pre-game one.	✓
H1 ₀	There is no statistical difference between the O1 objective pre-game and post-game test scores.	✓
H1 _A	The O1 objective post-game test score is statistically significantly higher than the pre-game one.	-
H2 ₀	There is no statistical difference between the O2 objective pre-game and post-game test scores.	✓
H2 _A	The O2 objective post-game test score is statistically significantly higher than the pre-game one.	-
H3 ₀	There is no statistical difference between the O3 objective pre-game and post-game test scores.	-
H3 _A	The O3 objective post-game test score is statistically significantly higher than the pre-game one.	✓
H4 ₀	There is no statistical difference between the O4 objective pre-game and post-game test scores.	-
H4 _A	The O4 objective post-game test score is statistically significantly higher than the pre-game one.	✓
H5 ₀	There is no statistical difference regarding the application play time between the experts and the non-experts.	✓
H5 _A	Non-experts show a statistically significantly higher application play time than experts.	-
H6 ₀	There is no statistical difference regarding the improvement between the experts and the non-experts.	✓
H6 _A	Non-experts show a statistically significantly higher improvement than experts.	-

and set an order in which these variations had to be played. Furthermore, to prevent any possible biases while they tried to solve the different game variations, we renamed the game variations Tufte Game, Few Game, and Mixed Game to *Game 1*, *Game 2*, and *Game 3*, respectively. After completing these requirements, the participants received a unique code necessary to continue with the user study. Once participants had finished the game variations, they were asked to take another visualization literacy test, the post-game test. Finally, the survey concluded with participants completing demographic and qualitative questions.

The pre- and post-game visualization literacy tests are identical to ensure the same difficulty level. Participants did not receive any feedback regarding the correctness of their answers. We also tracked the participants' interactions with Iguanodon, including the number of attempts and time taken to complete each game variation. The unique code generated for each participant was used as a matching key to connect the survey data with the tracked application data. We conducted a pilot study with four participants to refine textual descriptions, test the matching between the survey and game tracking data, estimate the study duration, and evaluate the functionality of the game application. We incorporated the feedback from the pilot study into our application, specifically regarding functionality and usability.

We recruited participants from among computer science students enrolled in introductory visualization courses at the Johannes Kepler University Linz. For a discussion of this choice of recruitment, see Section 5.3.

To evaluate the improvements in the proficiency in creating visualizations with our game application and to analyze the differences in the participants' visualization literacy experience, we formulated seven hypotheses, as listed in Table 3. H0 investigates the overall difference between the pre- and post-game test scores. A difference in the test scores indicates a change in proficiency in creating visualizations. We expected the overall test score to improve for the post-

game test (H0_A). Hypotheses H1 through H4 aim to investigate the differences in test scores related to four distinct objectives between the pre- and post-game tests. Based on our assumptions, we expected to observe an improvement in test scores across all four objectives for the post-game tests (H1_A to H4_A). H5 evaluates how long the game application play times differ between experts (participants with a lot of experience) and non-experts (with some or no experience). We assumed that non-experts are slower in solving the different game variations and therefore needed more time to finish the game (H5_A). For H6, we evaluated the difference in improvement between non-experts and experts by comparing pre- and post-game test scores, to assess whether non-expert users would exhibit greater improvements. (H6_A).

4.1 Visualization Literacy Test

To assess the visualization literacy of our study participants, we reviewed previously published works to identify a suitable test, as detailed in Section 2.1. However, many existing tests primarily focus on reading and interpreting tasks for different visualizations, such as the VLAT [5]. Other works, such as the data visualization literacy framework by Börner et al. [10], do provide some potential for evaluating the construction aspect of visualization literacy. Their framework combines a literature review and feedback from their Information Visualization course. Their construction assessment is based on creating visualizations that fulfill specific criteria in their topology. These visualizations are then evaluated by fellow students. However, the evaluation of the construction aspect does not apply to our needs (because there are no concrete questions and answers), and the topology does not cover all our objectives and actions.

We, therefore, opted to create our own visualization literacy test that focuses on evaluating the construction of a visualization. We tailored the visualization literacy test to our objectives and actions. We designed the test such that

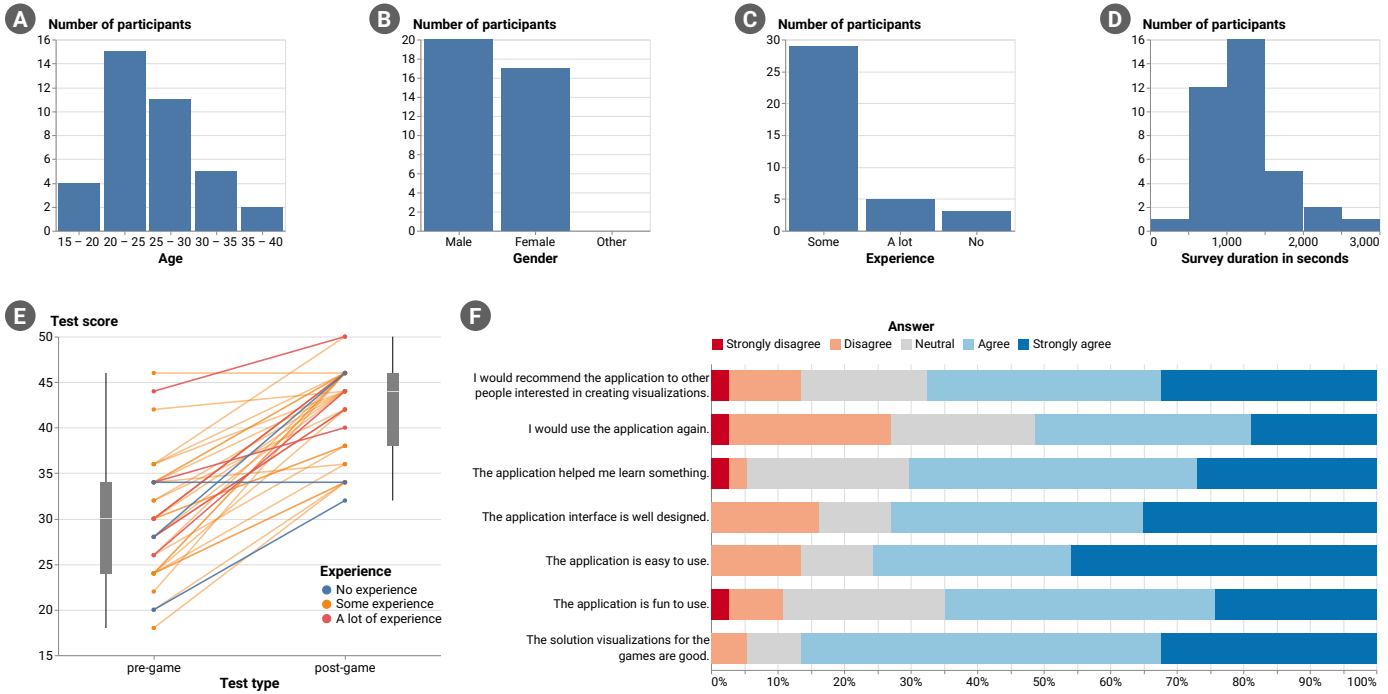


Fig. 3. Overview of the study participants’ age (A), gender (B), experience (C), and study duration (D) distributions. The test score improvements of the participants (E). The results for the seven Likert-scale questions are colored by their agreement level, indicating the high acceptance rate of our game application (F).

each answer option corresponds to an action in Iguanodon. We decided to test multiple objectives in each question to not overtax the participants with a long study duration.

The final result is a twelve-question multiple-choice test, with the first question about objective **O1** *Reduce overplotting*. This question includes a definition for overplotting, based on the work of Park et al. [44], and an example visualization where overplotting occurs. The remaining eleven questions cover all possible combinations of the four objectives, with at least two objectives addressed in each question. Each question follows the same format: the participant is presented with a visualization together with a list of possible answer options of how to improve the visualization. The visualizations and answer options used in the test match the ones in our game application.

We included the visualization literacy test in the supplementary material [45].

4.2 Analysis and Results

We started our survey analysis by evaluating the demographic data of the participants. Subsequently, we used the Mann-Whitney U test to determine whether the pre- and post-game or the non-expert and expert sets differ significantly. Lastly, we analyzed the results of the qualitative questions. There are different ways multiple-choice questions can be evaluated. For our test, we used a scoring system that adds one point for every correct answer and subtracts one point for every wrong answer. Using this method, the highest possible score would be 58 if all questions were answered correctly. We set the significance level of $\alpha = 0.05$, and used Bonferroni correction ($p_{corr} = n * p$, with p_{corr} being the corrected p -value and n the number

of tests). All p -values presented in the following section are corrected for multiple testing.

Thirty-seven participants completed the user study. Figure 3 (A–D) provides an overview of their characteristics. We had 20 male and 17 female participants aged between 17 and 39 ($M = 24.86, SD = 5.26$) with varying degrees of experience. The participants’ experience levels, ranging from *no experience* to *some experience* to *a lot of experience*, were self-reported. Most participants had some experience in creating visualizations. It took the participants from about 8 minutes to 43 minutes ($M = 20.95, SD = 8.12$) to finish the study.

4.2.1 Quantitative Analysis

We analyzed the statistical difference between the pre- and post-game test scores in hypotheses H0 – H4, as summarized in Figure 4. Figure 4 shows the results regarding hypotheses H0–H4. The plots show the mean pre- and post-test scores, with 1 meaning that all questions were answered correctly. We were able to confirm the difference of test scores before and after the use of Iguanodon, visible in Figure 3 (E), as statistically significant ($U = 1263.5, p < 0.0001$). Therefore we reject the null-hypothesis H0 and accept H0A. This verifies our goal to increase proficiency in creating visualizations with our game application. We further analyzed the improvement of the post-game test score regarding the objective **O1** *Reduce overplotting* ($U = 850, p = 0.2602$) and accept H10 and reject H1A. Subsequently, we took a look at the test scores concerning the objective **O2** *Reduce chartjunk* ($U = 869, p = 0.1618$), leading us to accept the null-hypothesis H20 and reject H2A. There is no significant improvement for the two objectives. For H3 we investigated




Fig. 4. The evaluation results of the seven different hypotheses with their mean value and their 95% confidence interval. H0 to H4 show the proportion of a perfect score for the pre- and post-game visualization literacy test. H5 and H6 show the difference between non-experts and experts for the playing time and test score difference, respectively.

the test results ($U = 1237.5, p < 0.0001$) about objective **O3** *Improve readability* and reject $H3_0$ and accept the alternative hypothesis $H3_A$. As for H4, the post-game test scores ($U = 1091, p < 0.0001$) are significantly higher for objective **O4** *Utilize color encoding properly* and we, therefore, accept $H4_A$ and reject $H4_0$.

The results of H0 to H4 indicate that our game application improves proficiency in creating visualizations. However, a more detailed analysis of the different objectives reveals that there is a significant improvement for only two out of the four objectives. We assume there are different reasons for this. As seen in Figure 4 H1, the questions related to objective **O1** *Reduce overplotting* were challenging to solve, with a mean of only approximately 9% correct answers in the pre-game test. The post-game test showed a small improvement, with a mean of around 16% of the correct answers for these questions. Therefore, it appears that there was not a significant improvement for this objective, possibly due to the difficulty of these questions or the participants did not think that the visualizations in the test suffered from overplotting. In contrast, the questions about objective **O2** *Reduce chartjunk* were less challenging. They led to a high pre-game test score with a mean of around 90% (Figure 4 H2), which makes it challenging to achieve further significant improvements.

Questions related to H3 and H4 in Figure 4 show a significant improvement in test scores for the objectives **O3** *Improve readability* and **O4** *Utilize color encoding properly*, respectively. A possible reason for this could be that the participants were not certain which actions should be applied to fulfill these objectives, and Iguanodon made them aware of it.

Additionally, we looked at the participants' playing

times regarding their experience level. We accept the null-hypothesis $H5_0$ ($U = 55, p = 4.7437$) and reject $H5_A$. The fact that there is no significant difference between the playing time of experts and non-experts could be that the game application does not impose a time limit, and therefore, there was no pressure for the participants to finish as fast as possible. We had to exclude one participant who was part of the expert group from calculating H5 because of a problem with saving their interactions on the server. Furthermore, the large confidence interval for the experts in Figure 4 H5 is attributable to the low number of participants with this experience level. Lastly, we investigate the improvement of the test scores with the distinction of experts and non-experts ($U = 81, p = 3.4399$), and we accept $H6_0$ and reject the alternative $H6_A$. Figures 4 H6 and 3 E show the improvements of the participants. It appears that experts and non-experts improved to the same degree. A reason for this could be the subjective judgment of the participants' experience level.

As a last step, we analyzed the logged interactions of the study participants for the different game variations. As a result, we could identify that the participants toggled the most actions during the first attempt of each game variation. Furthermore, we investigated the number of attempts the participants needed to solve the game variations. For the Tuft Game and Few Game, around 70% of the participants were able to solve them within the first two attempts. However, for the Mixed Game, that was only possible for about 40%, indicating greater difficulty due to the higher number of objectives and actions.

4.2.2 Qualitative Analysis

To further evaluate the usability and functionality, we added Likert-scale questions at the end of the questionnaire as well as two free text input fields. One input field was dedicated to collecting the aspects of the game application that participants enjoyed, while the other was intended for participants to express any aspects they disliked. In Figure 3 F, the visualized results of the Likert-scale questions show that the game application was generally well-received, with high agreement ratings (agree and strongly agree), e.g., for ease of use, learning, and enjoyment.

The positive feedback can be summarized as follows: The game application is well-designed, easy to use, and user-friendly, with clear instructions. Some participants liked the interactivity and the ability to see their actions' impact in the visualization immediately. The feedback via the objectives after each attempt was appreciated as well. With one participant commenting, "I liked that the feedback explained WHAT objective is not fulfilled." Another participant mentioned, "you get some feedback and instantly recognize which options you picked were wrong." Most participants found Iguanodon a fun way to learn and test visualization skills and a good starting point to learn and practice visualization concepts. One participant stated, "It could be a very good tool for playing around with these concepts so they can be learned in better and funnier ways."

The participants also provided valuable feedback about aspects they did not like and suggested ways to improve the game application. Some participants suggested adding textual explanations and reasoning behind certain actions.

One participant commented that they “would love some textual explanation. I see that this is hard to do, but for example, WHY is this one setting not good for visual clutter or similar questions?” Other participants wanted additional features, such as access to visualizations of previous attempts and different visualization types besides scatterplots. A participant mentioned, “Only scatter plots are used in the examples that I have seen. Maybe that’s only the case for me; if not, please add different kinds of plots.” A few participants experienced interface and layout issues, including scrolling problems, lack of responsiveness, and difficulty locating buttons. One participant said that “on smaller screens, the example chart uses a lot of space, and the section with the toggle gets very small, which makes the layout unclear.” This feedback helped us to identify issues where additional effort is needed to improve the overall usability of the game application.

5 DISCUSSION

5.1 Objectives and Actions

Iguanodon can be used to increase users’ proficiency in creating visualizations, but it is limited in its scope. The current version uses scatterplots only and includes a curated list of objectives and actions. However, the same approach can be used for different visualization types, objectives, and actions. Furthermore, this approach can be adapted to address specific topics, such as designing visualizations for people with color blindness or by defining objectives based on the media in which the visualization will be used. For example, a visualization intended for gray-scale print would require different objectives than one intended for an online platform. Another approach to creating effective visualizations for special contexts is establishing specific objectives tailored to the needs of an organization. Similarly, a tailored set of objectives and actions could be used for learning how to create visualizations for specific domains or to adhere to specific norms. For example, the DIN 461 norm [46], [47] is a standard used in engineering drawings to ensure consistency and accuracy. By incorporating this norm into our game applications, users could learn how to create effective visualizations that comply with this standard.

A large number of objectives and actions could increase the likelihood of contradicting design choice guidelines. An example of that would be the *no double or triple encoding guideline* by Tufte [31] and the recommendation by Franceneri [38] of double encoding for the color-blind viewers. In such a case, there may be no solution that fulfills all objectives. It may become necessary to prioritize objectives and actions based on what is most important for the intended audience. Textual explanations could be added to clarify why certain actions are preferable or that multiple solutions exist that fulfill the objectives.

5.2 Iguanodon Design

When designing the game, we thought about how to prevent users from applying a trial-and-error approach in which they randomly click actions in an attempt to achieve a high score. We mitigate this by requiring users to confirm each attempt before getting feedback on how they affected the objectives.

Furthermore, it is important to be aware that as the number of objectives and actions increases, the user interface can become overloaded and difficult to navigate. Therefore, it is essential to strike a balance between having enough objectives to achieve the desired learning outcome and having so many that it becomes difficult for the users to understand and navigate the interface.

5.3 User Study Participants

We decided to send the participation link to computer science student cohorts, taking into consideration their existing experience in creating visualizations. Most of these students have already developed a basic level of visualization literacy through conventional methods such as lectures and tutorials, making them an ideal group to evaluate our game application since we aimed to improve proficiency in creating visualizations by complementing the traditional approach with our game. In general, the process of increasing proficiency in creating visualizations should not significantly differ based on individuals’ backgrounds. Additionally, it is worth noting that using Iguanodon does not require prior computer science knowledge. However, the results of our user study are biased due to only having participants with a computer science background. To assess the effectiveness of Iguanodon for a general audience, an additional user study with participants with diverse backgrounds is needed.

5.4 Learning Effect

For our user study, we created a visualization literacy test that was presented to the participants before and after the game. We minimized potential learning effects by not providing participants with any feedback on their pre-game test performance, and we did not inform them that they needed to take the same test after playing the game. However, we encouraged participants to engage in active learning by using Iguanodon to find solutions to questions they remembered from the pre-game test. It is worth noting that we did not measure the long-term learning effect of the game. To do this, we would need to administer a third test several days or weeks after the game. This would have allowed us to assess how well participants retained the information they learned with our game application over an extended period of time.

5.5 Future Work

In future work, we plan to extend our game application by defining new game variations with different visualization types, such as line charts and bar charts. Also, since Iguanodon’s current version does not explain the objectives and their corresponding actions, as a next step, we want to add such explanations to give users a better understanding of why certain actions should be applied. Furthermore, these new explanations would allow us to provide feedback for contradicting objectives. Adding new game variations with contradicting objectives would highlight that the design choices are not always straightforward. Additionally, we want to improve the usability, mainly focusing on the feedback received from our study participants. Specifically, we recognize that smaller screens present a challenge as

they require excessive scrolling, making it more difficult for users to locate all the buttons within Iguanodon. Finally, Iguanodon can be extended so that, for example, teachers can quickly add new game variations with their own specifications without adapting the game application code.

6 CONCLUSION

Acquiring a high level of visualization literacy makes reading, interpreting, and constructing effective visualizations possible. These skills are becoming more important in our increasingly data-dependent world. In this paper, we present Iguanodon, a game application for improving proficiency in creating visualizations, superficially the proficiency of creating visualizations. Our approach builds upon the established design guidelines from previously published work and defines objectives that represent sub-tasks for the creation of effective visualizations. Iguanodon demonstrates its adaptability in creating different game variations based on different published works. Each objective consists of one or more design choices, which we defined as actions. We evaluated our game application by conducting a user study with 37 computer science students. To assess the visualization literacy of the participants, we designed our own visualization literacy test. The results of our user study indicate that our game application improves the proficiency in creating visualizations of the participants significantly.

ACKNOWLEDGMENTS

This work was supported by the Austrian Science Fund (FWF DFH 23-N), and the Austrian Research Promotion Agency (FFG 881844). Pro²Future is funded within the Austrian COMET Program under the auspices of the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology, the Austrian Federal Ministry for Digital and Economic Affairs, and of the States of Upper Austria and Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

REFERENCES

- [1] M. Solen, "Visualization Literacy in the Age of Big Data: Vital Skills for Modern Media Consumption," 2021.
- [2] Salesforce, Inc. (2023, May) Tableau: Business Intelligence and Analytics Software. [Online]. Available: <https://www.tableau.com/>
- [3] Microsoft. (2023, May) Data Visualization | Microsoft Power BI. [Online]. Available: <https://powerbi.microsoft.com/en-us/>
- [4] T. M. Connolly, E. A. Boyle, E. MacArthur, T. Hainey, and J. M. Boyle, "A systematic literature review of empirical evidence on computer games and serious games," *Computers & Education*, vol. 59, no. 2, pp. 661–686, 2012. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0360131512000619>
- [5] S. Lee, S.-H. Kim, and B. C. Kwon, "VLAT: Development of a Visualization Literacy Assessment Test," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 551–560, 2017.
- [6] K. Börner, A. Maltese, R. N. Balliet, and J. Heimlich, "Investigating aspects of data visualization literacy using 20 information visualizations and 273 science museum visitors," *Information Visualization*, vol. 15, no. 3, pp. 198–213, 2016. [Online]. Available: <https://doi.org/10.1177/1473871615594652>
- [7] J. Boy, R. A. Rensink, E. Bertini, and J.-D. Fekete, "A Principled Way of Assessing Visualization Literacy," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 1963–1972, 2014.
- [8] E. E. Firat, A. Joshi, and R. S. Laramée, "Interactive visualization literacy: The state-of-the-art," *Information Visualization*, vol. 21, no. 3, pp. 285–310, 2022. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/14738716221081831>
- [9] E. E. Firat, A. Joshi, and R. S. Laramée, "VisLitE: Visualization Literacy and Evaluation," *IEEE Computer Graphics and Applications*, vol. 42, no. 3, pp. 99–107, 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9790006/>
- [10] K. Börner, A. Bueckle, and M. Ginda, "Data visualization literacy: Definitions, conceptual frameworks, exercises, and assessments," *Proceedings of the National Academy of Sciences*, vol. 116, no. 6, pp. 1857–1864, 2019. [Online]. Available: <https://www.pnas.org/content/116/6/1857>
- [11] J. Mackinlay, P. Hanrahan, and C. Stolte, "Show Me: Automatic Presentation for Visual Analysis," *IEEE Transactions on Visualization and Computer Graphics (InfoVis '07)*, vol. 13, no. 6, pp. 1137–1144, 2007.
- [12] C. Stolte, D. Tang, and P. Hanrahan, "Polaris: A System for Query, Analysis, and Visualization of Multidimensional Relational Databases," *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 52–65, 2002.
- [13] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer, "Voyager: Exploratory Analysis Via Faceted Browsing of Visualization Recommendations," *IEEE Transactions on Visualization and Computer Graphics (InfoVis '15)*, vol. 22, no. 1, pp. 649–658, 2016.
- [14] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer, "Voyager 2: Augmenting Visual Analysis with Partial View Specifications," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 2017, pp. 2648–2659. [Online]. Available: <https://doi.org/10.1145/3025453.3025768>
- [15] H. Wu, D. Shi, N. Chen, Y. Shi, Z. Jin, and N. Cao, "VisAct: a visualization design system based on semantic actions," *Journal of Visualization*, vol. 23, no. 2, pp. 339–352, 2020. [Online]. Available: <http://link.springer.com/10.1007/s12650-019-00617-x>
- [16] K. Hu, M. A. Bakker, S. Li, T. Kraska, and C. Hidalgo, "VizML: A Machine Learning Approach to Visualization Recommendation," in *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '19)*. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–12. [Online]. Available: <https://doi.org/10.1145/3290605.3300358>
- [17] V. Dibia and Ç. Demiralp, "Data2Vis: Automatic Generation of Data Visualizations Using Sequence-to-Sequence Recurrent Neural Networks," *IEEE Computer Graphics and Applications*, vol. 39, no. 5, pp. 33–46, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8744242/>
- [18] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer, "Vega-Lite: A Grammar of Interactive Graphics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 341–350, 2017.
- [19] Y. Luo, X. Qin, N. Tang, and G. Li, "DeepEye: Towards Automatic Data Visualization," in *Proceedings of the IEEE International Conference on Data Engineering (ICDE '18)*. Paris: IEEE, 2018, pp. 101–112. [Online]. Available: <https://ieeexplore.ieee.org/document/8509240/>
- [20] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer, "Formalizing Visualization Design Knowledge as Constraints: Actionable and Extensible Models in Draco," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 438–448, 2019.
- [21] H. Li, Y. Wang, S. Zhang, Y. Song, and H. Qu, "KG4Vis: A Knowledge Graph-Based Approach for Visualization Recommendation," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 195–205, 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9552844/>
- [22] Q. Chen, F. Sun, X. Xu, Z. Chen, J. Wang, and N. Cao, "Vi-ZLint: A Linter and Fixer Framework for Data Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 206–216, 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9552878/>
- [23] A. K. Hopkins, M. Correll, and A. Satyanarayan, "VisuLint: Sketchy In Situ Annotations of Chart Construction Errors," in *Computer Graphics Forum (Proc. EuroVis)*, ser. 3, vol. 39, 2020, pp. 219–228. [Online]. Available: <http://vis.csail.mit.edu/pubs/visuallint>

- [24] B. Alper, N. H. Riche, F. Chevalier, J. Boy, and M. Sezgin, "Visualization Literacy at Elementary School," in *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI '17)*, ser. CHI '17. ACM, 2017, pp. 5485–5497. [Online]. Available: <https://doi.org/10.1145/3025453.3025877>
- [25] F. Bishop, J. Zagermann, U. Pfeil, G. Sanderson, H. Reiterer, and U. Hinrichs, "Construct-A-Vis: Exploring the Free-Form Visualization Processes of Children," *IEEE Transactions on Visualization and Computer Graphics*, pp. 451–460, 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/8807271/>
- [26] J. Gäbler, C. Winkler, N. Lengyel, W. Aigner, C. Stoiber, G. Wallner, and S. Kriglstein, "Diagram Safari: A Visualization Literacy Game for Young Children," in *Proceedings of the Symposium on Computer-Human Interaction in Play Companion Extended Abstracts (CHI PLAY '19)*. Barcelona Spain: ACM, 2019, pp. 389–396. [Online]. Available: <https://dl.acm.org/doi/10.1145/3341215.3356283>
- [27] E. Huynh, A. Nyhout, P. Ganea, and F. Chevalier, "Designing Narrative-Focused Role-Playing Games for Visualization Literacy in Young Children," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 924–934, 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9222251/>
- [28] S. He and E. Adar, "VizITCards: A Card-Based Toolkit for Infovis Design Education," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 561–570, 2017.
- [29] L. Amabili, K. Gupta, and R. G. Raidou, "A Taxonomy-Driven Model for Designing Educational Games in Visualization," *IEEE Computer Graphics and Applications*, vol. 41, no. 6, pp. 71–79, 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9556564/>
- [30] D. Rees and R. S. Laramée, "A Survey of Information Visualization Books," *Computer Graphics Forum*, vol. 38, no. 1, pp. 610–646, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/10.1111/cgf.13595>
- [31] E. R. Tufte, *The Visual Display of Quantitative Information*, 2nd ed. Graphics Press, 2001.
- [32] S. Few, *Show Me the Numbers: Designing Tables and Graphs to Enlighten*, 2nd ed. Analytics Press, 2012.
- [33] G. Ellis and A. Dix, "A Taxonomy of Clutter Reduction for Information Visualisation," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1216–1223, 2007.
- [34] J. Camões, *Data at work: best practices for creating effective charts and information graphics in Microsoft Excel*. New Riders, 2016.
- [35] Y. Holtz and C. Healy. (2023, May) From data to Viz. [Online]. Available: <https://www.data-to-viz.com/>
- [36] A. Kirk, *Data Visualisation: A Handbook for Data Driven Design*. Sage Publications Ltd, 2016.
- [37] D. A. Keim, "Visual Techniques for Exploring Databases," in *Conference on Knowledge Discovery in Databases (KDD'97)*, 1997.
- [38] S. L. Franconeri, L. M. Padilla, P. Shah, J. M. Zacks, and J. Hullman, "The Science of Visual Data Communication: What Works," *Psychological Science in the Public Interest*, vol. 22, no. 3, pp. 110–161, 2021. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/152910062111051956>
- [39] U. C. Bureau. (2023, May) Data Visualization Standards. [Online]. Available: <https://xdgov.github.io/data-design-standards/>
- [40] Wikimedia Foundation, Inc. (2023, May) Mastermind (board game). [Online]. Available: [https://en.wikipedia.org/w/index.php?title=Mastermind_\(board_game\)&oldid=1143111323](https://en.wikipedia.org/w/index.php?title=Mastermind_(board_game)&oldid=1143111323)
- [41] American Statistical Association (ASA). (2023, May) Before 1983 ASA Data Exposition dataset. [Online]. Available: <https://community.amstat.org/jointscsg-section/dataexpo/dataexpobefore1993>
- [42] UW Interactive Data Lab. (2023, May) Arquero. [Online]. Available: <https://github.com/uwdata/arquero>
- [43] LimeSurvey GmbH. (2023, May) LimeSurvey — Free Online Survey Tool. [Online]. Available: <https://www.limesurvey.org/en/>
- [44] D. Park, S.-H. Kim, and N. Elmqvist, "Gatherplot: A Non-Overlapping Scatterplot," 2023. [Online]. Available: <https://arxiv.org/abs/2301.10843>
- [45] P. Adelberger, O. Lesota, K. Eckelt, M. Schedl, and M. Streit. (2023, May) Supplementary Material for Iguanodon. [Online]. Available: <https://osf.io/j95ps/>
- [46] DIN Deutsches Institut für Normung e. V. (2023, May) DIN 461. [Online]. Available: <https://www.din.de/en/getting-involved/standards-committees/natg/publications/wdc-beuth: din21:710844>


- [47] Wikimedia Foundation, Inc. (2023, May) DIN 461. [Online]. Available: https://de.wikipedia.org/w/index.php?title=DIN_461&oldid=231436431



Patrick Adelberger is a PhD student at the Institute of Computer Graphics, Johannes Kepler University (JKU) Linz, Austria. His research interests include visualization literacy, visualization for biomedical data, and visual analytics. He received his Diplomingenieur (MSc) in Medical Informatics from TU Wien, Vienna, Austria.



Oleg Lesota is a PhD student at the Institute of Computational Perception, Multimedia Mining and Search group, Johannes Kepler University and Human-centered Artificial Intelligence group at the Linz Institute of Technology AI Lab Linz, Austria. His research interests revolve around recommender systems, bias and fairness, visualization literacy, and visual analytics.



Klaus Eckelt is a PhD student at the Visual Data Science Lab at Johannes Kepler University Linz, Austria. His primary research focus is on facilitating cohort analysis in the biomedical domain, utilizing data visualization, visual analytics, as well as statistical and machine learning methods. For more information, see <https://eckelt.info>.



Markus Schedl is a Full Professor at Johannes Kepler University (JKU) Linz, Austria, affiliated with the Institute of Computational Perception, where he leads the Multimedia Mining and Search (MMS) group. In addition, he is head of the Human-centered Artificial Intelligence (HCAI) group at the Linz Institute of Technology (LIT) AI Lab. His areas of expertise include recommender systems, user modeling, information retrieval, machine learning, natural language processing, multimedia, data analysis, and web mining. For more information, see <http://mschedl.eu>.



Marc Streit is a Full Professor at the Johannes Kepler University Linz in Austria where he leads the JKU Visual Data Science Lab (<https://jku-vds-lab.at/>). He finished his PhD at the Graz University of Technology in 2011. His scientific areas of interest include visualization, visual analytics, and explainable AI. He won multiple best paper and honorable mention awards at major conferences in the field. Marc is also co-founder and CEO of datavisyn, a spin-off company that develops data visualization solutions for biomedical R&D. For more information, see <http://marc-streit.com>.