

# THE IDEALVIS ADAPTATION ENGINE

---

## D15

### IDEALVIs Consortium

<http://idealvis.inspirecenter.org/>



**European Union**

European Regional  
Development Fund



**Republic of Cyprus**



**Structural Funds**  
of the European Union in Cyprus



**ΙΔΡΥΜΑ  
ΕΡΕΥΝΑΣ ΚΑΙ  
ΚΑΙΝΟΤΟΜΙΑΣ**

# Executive Summary

Analytical processing is often a complex and cumbersome process, turning data to insights and insights to better and faster business decisions. Data visualization is among the most powerful tools in the analyst's arsenal, shedding light into established or previously unseen patterns. However, its success is frequently an interplay between the data analyst's experience and ability to quickly understand and interpret information. Even though business analytics tools have made a significant progress to deliver immersive information visualization environments for improving users' efficiency and effectiveness, they still do not consider incorporating individual differences in the core process that influences the visualization structure, encoding and readability. This paper leverages the users' individual differences to deliver a human-centred by-design adaptation engine for business users aiming to improve their comprehension of data visualizations, thus leading to improved accuracy and time-to-action efficiency. The proposed adaptation mechanism is evaluated using 45 professional analysts from multiple industry sectors. The results suggest that individual differences can play an important role in the adaptation process of data visualizations enhancing analysts' comprehensibility and decision making.

# Human-centered Information Visualization Adaptation Engine

Analytical processing is often a complex and cumbersome process, turning data to insights and insights to better and faster business decisions. Data visualization is among the most powerful tools in the analyst’s arsenal, shedding light into established or previously unseen patterns. However, its success is frequently an interplay between the data analyst’s experience and ability to quickly understand and interpret information. Even though business analytics tools have made a significant progress to deliver immersive information visualization environments for improving users’ efficiency and effectiveness, they still do not consider incorporating individual differences in the core process that influences the visualization structure, encoding and readability. This paper leverages the users’ individual differences to deliver a human-centred by-design adaptation engine for business users aiming to improve their comprehension of data visualizations, thus leading to improved accuracy and time-to-action efficiency. The proposed adaptation mechanism is evaluated using 45 professional analysts from multiple industry sectors. The results suggest that individual differences can play an important role in the adaptation process of data visualizations enhancing analysts’ comprehensibility and decision making.

CCS Concepts: • **Human-centered computing** → *Visualization toolkits*; **Visual analytics**; **Information visualization**; **User models**; **User centered design**; **Visualization techniques**; • **Information systems** → **Personalization**.

Additional Key Words and Phrases: information visualization, adaptation, personalization, user modeling, individual differences, rule based adaptation

## ACM Reference Format:

. . Human-centered Information Visualization Adaptation Engine. In . ACM, New York, NY, USA, 19 pages.

## 1 INTRODUCTION

The last decade has been a witness to phenomenal growth in the volume of information and data science, revolutionizing many industry domains. With such volumes of data being generated, companies that want to stay competitive in today’s data driven market quickly adopt business intelligence and analytics (BI&A) software. These software provide a collection of tools for handling, integrating, cleaning data and performing analysis of structured and unstructured data for the purpose of delivering actionable information to business decision makers [20]. Results of the performed analyses are usually presented using interactive data visualizations/data tables that are embedded within reports and dashboards that the user can utilize for performing further exploration, collaboration and acquisition of useful insights. Moreover, BI&A platforms are slowly adopting techniques that enable non-expert-analysts e.g., business users to utilize such tools. One of these techniques is self-service-analytics, which is described as “all the facilities of a BI environment that enable users to become more self-reliant and less dependent on the IT organization”; facilities may include easier access to data and simpler, customizable user-interfaces amongst others [6]. Despite the assistance provided to the user through the various techniques and facilities, the increasing numbers of data sources and the huge variety of data formats, can result in complications that tend to disorient the user that has to interact with rich business data visualizations with an abundance of features in order to reach a decision. Although, BI&A platforms provide vast amounts of customizable data visualizations they tend to follow a one-size-fits-all approach, where the rendered data visualizations are solely based

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© Association for Computing Machinery.

Manuscript submitted to ACM

53 on the selected dataset metadata or the current analysis task, not considering the user’s requirements or individual  
54 differences.

55 Research on individual differences in visualization is progressively growing, showing that interaction with a data  
56 visualization can be affected by an individual user’s cognitive abilities e.g., visual working memory [29], cognitive styles  
57 e.g., field dependent [27], personality factors e.g., extraversion [14] and expertise/experience [17, 18]. It is evident that  
58 the impact of individual differences on data visualizations needs to be further explored in the context of the business  
59 domain, since related research was rarely applied to such context. Subsequently, appropriate adaptive data visualization  
60 systems need to be developed, which consider the unique needs and characteristics of the business analyst user. This paper  
61 presents a novel human-centered adaptation engine, which aims to enhance the comprehensibility of data visualizations  
62 thus improving accuracy and time-to-action efficiency. The adaptation engine adopts an ensemble learning framework  
63 consisting of two phases: (i) a fuzzy rule generation procedure, which generates fuzzy (weak learner) rules based on the  
64 impact of cognitive factors on the understandability of data visualizations; and (ii) the classification procedure, which  
65 combines the rules to produce the adapted content.

66 The adaptation engine was evaluated, and the delivered adaptation was found to be effective in improving the user’s  
67 (i) performance (i.e., time taken to address an analysis task) and accuracy (i.e., correctness of analysis task response), as  
68 well as (ii) the perceived user experience and system usability scores. For performance users were faster by an average  
69 of 8.1 seconds when adaptation was enabled. Moreover, analysis task accuracy scores revealed that 62% of users were  
70 more accurate when responding to analysis tasks for which adaptation was enabled. Finally, adaptation impacted the  
71 users’ perceived user experience score with an increase of 9%, and the reported system usability score with an increase  
72 of 1.8%.

73 The rest of this paper is structured as follows: Section 2 presents the system model, consisting of the user model,  
74 analysis tasks, data and data visualizations, which form the core adaptation engine inputs. Next, Section 3 introduces the  
75 adaptation engine and its two phases: adaptation rules and adaptation process. The rule extraction process is further  
76 described in Section 4.6. A user study, consisting of 45 business analysts was performed to evaluate the adaptation  
77 engine’s impact on user’s efficiency and effectiveness. Section 5 presents the user study setup and discusses its results.  
78 Then, Section 8 presents related work on data visualization adaptation based on individual differences. Finally, Section 9  
79 concludes the paper.

## 80 2 SYSTEM MODEL

81 In this section we provide an overview of the system model. A table of respective symbols is provided in Table 1. The  
82 system has a set of data analyst users  $U$  and for each user  $\{u_1, u_2, \dots, u_N\}$  the system maintains a user model  $um(u_i)$ ,  
83 which stores the user’s characteristics (e.g., demographics, psychometric indicators). We assume that as part of their  
84 daily operations, each user  $u_i$  is assigned a subset of the organization’s data analysis tasks  $T = \{t_1, t_2, \dots, t_N\}$  that need  
85 to be addressed through appropriate explorations. We also assume that a task  $t_i$  will require the construction of some  
86 query  $q$  to retrieve some data  $d$ , which will be visualized appropriately by the user’s interface.

### 87 2.1 User Model

88 The user model of user  $um(u_i)$  is a set of triplets of the form  $(ct, ch, val)$ , where  $ct$  represents the user model item  
89 category (e.g.,  $d$ =demographics,  $p$ =psychometric characteristics),  $ch$  represents an actual characteristic that belongs to  
90 the triplet’s category (e.g., age), and  $val$  which represents the respective value for the triplet’s characteristic (e.g., 35 for  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104

Table 1. Table of Symbols

Symbol	Description
$\mathcal{U}$	Set of all users ( $\{u_1, u_2, \dots, u_N\}$ )
$um(u_i)$	The User Model for a given user $u_i$
$T$	Set of all system analysis tasks ( $\{t_1, t_2, \dots, t_N\}$ )
$V$	Set of all data visualization types ( $\{bar, column, line, radar, pie, table\}$ )
$VE$	Set of all modifiable visual elements ( $\{hgl, cpt1, cpt2, dt, esiz, prox, dl\}$ )
$q$	A user-defined query produced via the Analysis Wizard interface
$d$	Dataset that resulted from a user query $q$
$ARVT$	Adaptation rule set used for selecting the data visualization type ( $v_i \in V$ )
$ARVE$	Adaptation rule set used for selecting visual element modifications ( $ve_i \in VE$ )

age). An example of a user model for user  $u_i$  can be  $um(u_i) = \{(d, age, 35), (p, wm, low), (p, fdi, fd)\}$ , denoting that the user  $u_i$  has an age of 35, a low Working Memory and is classified as field-dependent.

## 2.2 Tasks

A task ( $t_i \in T$ ) represents a business question, such as "Identify if the glass bottles pack type is growing in terms of sales value in August of 2021 compared to June 2021.", that needs to be explored and addressed. A task is a tuple  $t_i = (text, ty_i)$  consisting of: (i) the narrative ( $text$ ); and (ii) the task's type ( $ty_i \in TY$ ), as summarised in Table 2 (following the work of Amar et al. [1] which presented a set of low-level analytical tasks that largely capture people's activities while employing information visualization tools). An example of a task  $t_i$  is ( $text =$  "Identify if all Soft Drinks sales were affected by seasonality in 2019",  $fan$ ), where the type of the task  $ty_i$  is  $fan = FindAnomaly$ ).

Table 2. Available System Task Types  $TY$ 

Task Type $ty_i$	Task Name	Task Descriptions [1]
$rol$	Retrieve Value	Given a set of specific cases, find attributes of those cases.
$cdv$	Compute Derived Value	Given a set of data cases, compute an aggregate numeric representation of those data cases.
$fan$	Find Anomalies	Identify any anomalies within a given set of data cases with respect to a given relationship or expectation, e.g., statistical outliers.
$cor$	Correlate	Given a set of data cases and two attributes, determine useful relationships between the values of those attributes.
$comp$	Simple Comparison	Simple data value comparison e.g., finding the lowest or the highest value in the dataset. (not in [1])

## 2.3 Data

The system employs an information retrieval engine that can support the users' data explorations. In the context of this work, we assume that the system maintains a number of datasets. The datasets can be accessed via a Query Engine, which can specify queries in a supported language, such as SQL. Following the example of the previous sections, to address the task  $t_i$ , the system can construct query  $q = "RetrieveMonth, SalesValue from SalesDataset"$ .

Table 3. Table of Adaptive Elements  $VE$  and their Applicability to Data Visualization Types

Element $ve_i$	Description	Bar Chart	Column Chart	Line Chart	Radar Chart	Pie Chart	Data Table
<i>hgl</i>	Enables horizontal and vertical grid lines	×	×	×	×		
<i>cpt1</i>	Switches to color palette 1 (duller colors compared to <i>cpt2</i> )	×	×	×	×	×	
<i>cpt2</i>	Switches to color palette 2 (brighter colors compared to <i>cpt1</i> )	×	×	×	×	×	
<i>dt</i>	Enables dark background and white text (dark theme)	×	×	×	×	×	×
<i>esiz</i>	Changes the default size of primary elements (bars, columns, lines)	×	×	×	×		
<i>prox</i>	Changes the default proximity between primary elements (bars and columns)	×	×				
<i>dl</i>	Displays data values on top of elements (e.g., above bars)	×	×	×	×	×	

## 2.4 Data Visualizations

The system maintains a Data Visualization Engine that is responsible for rendering all system’s data visualizations. All data visualizations  $V$ , have a set of applicable visual elements (i.e., visual modifications) ( $ve_i \in VE$ ) which can be enabled by the adaptation engine for delivering the desired adaptation/personalization. Table 3 lists the set of visual elements  $VE$  and their applicability on each of the available data visualization types  $V$ . In order to render a data visualization, the Data Visualization Engine requires: (i) the data returned by the Query Engine; (ii) the type of data visualization  $v_i$  to render; and (iii) the set of visual element modifications  $ve_i$  as selected by the Adaptation Engine.

## 3 ADAPTATION ENGINE

The adaptation engine is responsible for delivering the best fit data visualization to the analyst user by utilizing the user’s model  $um(u_i)$ , an analysis task  $t_j$ ; and two sets of adaptation rules for selecting: (i) the best fit data visualization type according to the user model; and (ii) a set of visual elements modifications, which further personalise the data visualization.

This section presents the internal mechanisms of the proposed adaptation engine, describing how the adaptation rules are utilized to deliver best-fit visualizations. Section 4 provides insight on how the rules were constructed.

### 3.1 Adaptation Rules

The adaptation rules represent *if-then* rules that are triggered during the adaptation process. They are stored in two sets: (i) rules for adapting visualization types ( $ARVT$ ); and (ii) rules for adapting visualization elements ( $ARVE$ ). Example rules from the  $ARVT$  and  $ARVE$  rule sets can be seen in Table 4 and Table 5 respectively.

The rules consist of several metadata attributes that allow them to be triggered according to the user, task and data characteristics. For example, the *Factor* and *Level* columns in both rule sets refer to human factors (e.g., Working Memory) and the level for that human factor (e.g., High) respectively. The rule sets are also augmented with several other metadata attributes (e.g., TaskType, TS and UniDim columns in the  $ARVT$ ) to provide further refinements for the analysis task type (e.g., Find Anomaly corresponding to task type  $ty_i = fan$ ), whether the analysis task processes

Table 4. Example Adaptation Rules for Visualization Type (ARVT)

Factor	Level	TaskType	TS	UniDim	BAR	RADAR	COLUMN	LINE	PIE	TABLE
WM	High	Find Anomaly	False	True	4 %	0 %	96 %	0 %	0 %	0 %
WM	High	Find Anomaly	True	True	0 %	92 %	0 %	4 %	0 %	4 %

Table 5. Example Adaptation Rules for Visual Elements (ARVE)

Factor	Level	ChartType	Element	ENABLE	DISABLE
WM	High	Bar	DARK_MODE	60 %	40 %
WM	High	Line	DARK_MODE	32 %	68 %

time-series data, and if the data feature one or more dimensions. Additionally, the *ARVE* rule set contains the *ChartType* and *Element* columns, which refer to the chart type and visual element modification that a specific rule applies.

The above-mentioned columns are used for filtering each of the rule sets prior to adaptation. For instance the adaptation engine uses Factor, Level, TaskType, TS and UniDim columns of the *ARVT* rule set for filtering down the rules according to user model attributes and analysis task metadata. The filtered rules are then used for voting the best visualization type. Similarly, the adaptation engine uses the Factor, Level, ChartType and Element columns of the *ARVE* rule set for filtering down the rules. It then uses the filtered rules for voting whether or not a specific visual element adaptation (e.g., dark theme *dt*) should be applied on the selected data visualization.

The selected representation of the rules serves two purposes: (i) to facilitate efficient execution of the rules; and (ii) to improve the explainability of the inference process carried out by the system. A simplified example of the first rule in Table 4 is presented below: IF Factor="WM" and Level="High" and TaskType="Find Anomaly" and TS="False" and UniDim="True" THEN VisualizationType=Column=0.96, Bar=0.04.

The rule is interpreted as follows: if the current user has high working memory, the active task is Finding Anomaly, the data are not time series data and they contain one dimension, then *vote* for the column visualization type with 96% and for the bar visualization type with 4%. As it will be described in the next Section, the ensemble of all rules, considering all aspects of the user model, yields the final result, i.e., the visualization type with the highest weighted vote. The same voting approach is used for rules in the *ARVE* rule set, instead for that rule set the engine votes whether or not to enable a specific visual element modification.

### 3.2 Adaptation Procedure

The complete adaptation procedure is illustrated in Figure 1. The procedure is initiated when a user attempts to retrieve data for a task  $t_i$  (step 1a) using a query  $q$ . Before any adaptation takes place, the Adaptation Engine forwards the query to the Query Engine (step 1b). From there the Query Engine is responsible for processing  $q$  (step 1c) and returning the required data  $d$  for the analysis (step 1d) back to the Query Engine. Subsequently, the Query Engine returns the data (step 1e) back to the Adaptation Engine, which are combined with the task's metadata  $tm_i$  (step 1f).

The Adaptation Engine can take two possible routes of execution: (a) perform adaptation (step 2a); or (b) skip adaptation (step 3). This is decided according to the application requirements (e.g., a task that has a pre-selected data visualization type) or user study setup (e.g., testing non-personalized vs. personalized content). If no personalization is required, the adaptation process is skipped, and the data (step 4) and pre-selected data visualization type (step 3) are

261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312

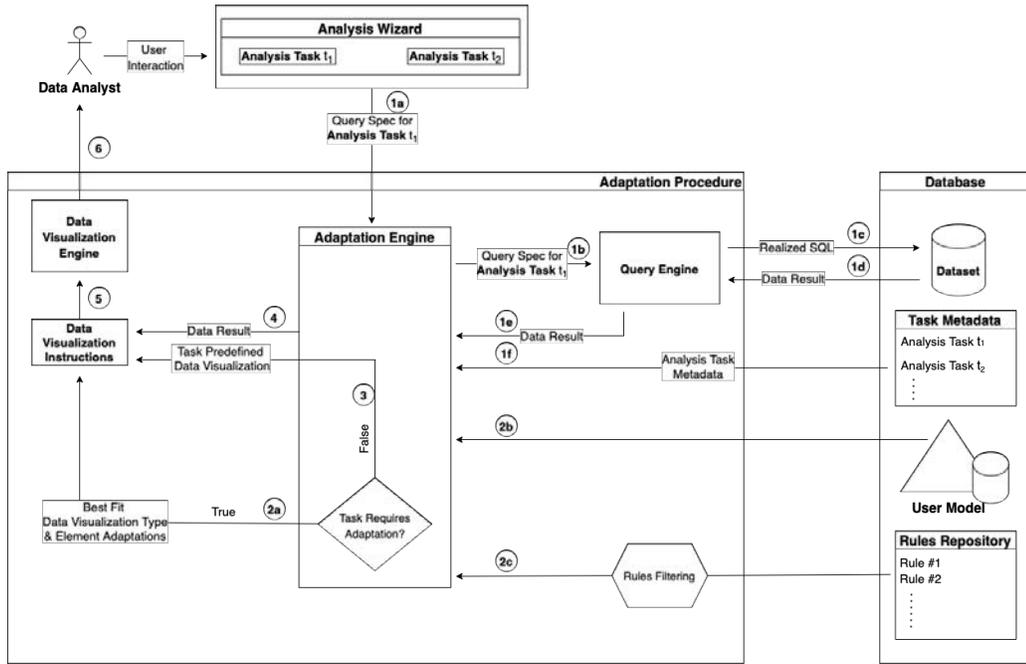


Fig. 1. Adaptation Process

forwarded to the Data Visualization Engine (step 5) that is responsible for rendering the output to the user (step 6). Alternatively, if adaptation is required, the user's model  $um(u_i)$  (step 2b) and task metadata  $tm_i$  are used for filtering the rule sets and retrieving all matching rules (step 2c). The aggregation of all rule votes decides first the best fit data visualization type and then accordingly the visual element modifications. Finally, the "winner" data visualization type and its visual element modifications/adaptations are sent to the Data Visualization Engine (step 5) that is responsible for rendering the adapted output to the user (step 6).

The adaptation procedure can be conceptually illustrated as the following function:

$$ap(q, t_i, u_i) = (v_i, d, VE) \quad (1)$$

The adaptation procedure function receives a query  $q$  for a task  $t_i$  by user  $u_i$ , and returns the desired data visualization instructions in the form of a triplet that contains the best fit data visualization ( $v_i$ ), the queried data  $d$  and any visual element modifications ( $ve_i \in VE$ ). The output of this function is used as input to the Data Visualization Engine. The logic behind this function is further detailed using pseudo-code in Algorithm 1.

The adaptation procedure algorithm 1 starts by retrieving the queried data  $d$  for query  $q$  (line #1) and the task metadata  $tm_i$  (line #2). If the task has a predefined data visualization (line #3) the algorithm stops the execution and returns the pre-selected data visualization type for the task, the queried data and the default set of visual element modifications  $\{\}$  (line #4). Otherwise, adaptation is initiated by first retrieving the user's model (line #6) and subsequently identifying the best fit data visualization type. This is done by first filtering down the ARVT rule set (line #7) according to the user's user model and task metadata. The resulting set of ARVT rules is forwarded to the voter function (bagging

**Algorithm 1** The Adaptation Procedure Algorithm  $ap(q, t_i, u_i)$ 


---

```

313 1:  $d \leftarrow queryEngine(q)$ 
314 2:  $tm \leftarrow getTaskMetadata(t_i)$ 
315 3: if  $tm.predefinedVisualization$  is not NULL then
316 4:   return  $(tm.predefinedVisualization, d, \{\})$ 
317 5: end if
318 6:  $um \leftarrow um(u_i)$ 
319 7:  $ARVT \leftarrow filterARVT(um, t_i)$ 
320 8:  $bestVisType \leftarrow visTypeVoter(ARVT)$ 
321 9:  $selectedElements \leftarrow \emptyset$ 
322 10: for each  $el_i \in elements$  do
323 11:    $ARVE \leftarrow filterARVE(um, bestVisType, el_i)$ 
324 12:    $enableElement \leftarrow visElementVoter(ARVE)$ 
325 13:   if  $enableElement$  then
326 14:      $selectedElements \leftarrow selectedElements \cup \{el_i\}$ 
327 15:   end if
328 16: end for
329 17: return  $(bestVisType, d, selectedElements)$ 

```

---

ensemble aggregation) that aggregates the votes for each of the available visualization types and returns the *winner* data visualization (line #8).

The algorithm then selects a number of visual element modifications that are specific to the winner data visualization. Initially an empty set called `selectedElements` is initiated (line #9) for keeping element modifications that are to be enabled on the winner data visualization. Moreover, at (line #10) a for-loop goes through all possible visual element modifications (*VE*) supported by the visualization engine and filters rules from the *ARVE* rule set using the user's user model, the winner data visualization type and the visual element that is to be voted. The filtered set of rules is forwarded to the voter function (line #12) that aggregates the votes and decides to enable or disable the current visual element modification  $ve_i$ . The result of the vote is added in the `selectedElements` set (line #14). Finally, once all elements are evaluated the winner data visualization type, the queried data and the `selectedElements` set are returned to the Visualization Engine (line #17).

#### 4 ADAPTATION RULES EXTRACTION STUDY

Adaptation rules are used for selecting the best fit data visualization and modifier visual elements in accordance to the business user's characteristics. This section describes how the the adaptation rule-sets and user models were extracted through a user study that analyzed the interactions (performance and accuracy) of 60 business data analyst participants with different data visualizations. The study required participants to engage with various visual analysis tasks of varying complexity, with each task constructed using different visualization conditions, recording their performance and accuracy. The study enabled us to assess the impact of individual differences and human factors on the performance and accuracy when solving visual analysis tasks, influenced by different visualization types, and by different visual elements i.e., visual modifications.

## 4.1 Participants

The study recruited 60 business data analyst participants from two industry organizations (RAI Consultants LTD<sup>1</sup> and KPMG Cyprus<sup>2</sup>) that had on average at least 2 years of experience in the field of data analytics. An important participation precondition was that part of the participant's daily job responsibilities was the interaction with data visualizations. Moreover, the sample was made up of analysts working on diverse industry fields such as Retail, Marketing, Advisory Services, Audit and Risk Assessment, consisting of 30 Male and 30 Female participants, with their ages ranging from 24 to 57 (mean age  $33.9 \pm 7.8$ ). All participants' had varying expertise levels, including directors and managers, executive-, senior- and junior-analysts.

*4.1.1 User Modeling.* The user model, which includes human factors, demographics, experience and expertise, of all users were extracted using psychometric tests and questionnaires based on our previous works in [2, 31]. In particular, the user model is comprised of the Field Dependent-Independent (FDI) cognitive style construct, several cognitive processing abilities and traditional characteristics. The FDI values were captured using a virtual version of the Group Embedded Figures Test [34]. The cognitive abilities include: (i) two Stroop-like tests used for capturing the control of attention and speed of processing; and (ii) a visual working memory test similar to the one in [9]. Furthermore, user model includes other factors, such as: (i) personality factors, extracted using the Eysenck Personality Questionnaire (EPQR-S) English version [10]; (ii) emotional regulation factors, extracted using the Emotional Regulation Questionnaire (ERQ) [15]; (iii) decision making factors, extracted using the Decision Making Style Questionnaire [23]; and (iv) problem solving factors, extracted using Parker's Problem-Solving Style Questionnaire (PSSQ) [11]. Finally, the user modeling process also records demographic information for each participant including their gender, age, and their educational status, and experience and expertise characteristics.

## 4.2 Dataset

To cater for the diverse user analysts' expertise, a synthetic sales dataset was constructed, with concepts that are easy for comprehension and analysis. More specifically, a synthetic dataset of comic book sales was constructed, with typical dimensions, such as time, product characteristics and location characteristics, and a few distributive (e.g., quantity, price) and algebraic measures (e.g., average price, weighted price).

## 4.3 Data Visualization Types

According to our previous findings [2] we decided to use the data visualizations seen in Table 6 for this user study. These visualizations represent the most frequently used data visualizations in the business domain. While the Radar Chart is less widely used in the domain of interest when compared to the rest of the selected charts, we decided to include it since participants with higher visual working memory may have higher preference for this type of chart over bar charts [29].

## 4.4 Data Visualization Tasks

Using the synthetic comic book sales dataset, we further produced 160 visual exploration tasks. To facilitate our description, consider the task example presented in Figure 2. A task consists of a business question that a participant needs to address, accompanied by a non-interactive data visualization, which is not interactive. Finally, at the bottom of

<sup>1</sup><https://www.rai.com.cy/>

<sup>2</sup><https://home.kpmg/cy/>

Table 6. Available System Data Visualizations  $V$ 

Data Visualization Type $v_i$	Visualization Name
<i>bar</i>	Bar Chart
<i>line</i>	Line Chart
<i>pie</i>	Pie Chart
<i>column</i>	Column Chart
<i>table</i>	Data Table
<i>radar</i>	Radar Chart

the task, the user can provide their response and submit their answer. In the background, the system records the time it took for the participant to respond (in milliseconds) and assess whether the response of the participant was correct.

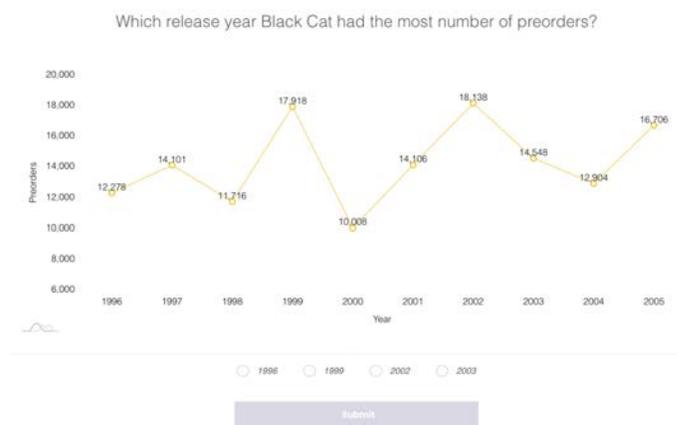


Fig. 2. User Interface for Addressing a Data Analysis Task

The tasks were organized into four experiments: (i) *Chart Type* experiment; (ii) *Task Complexity* experiment; (iii) *Dimensionality* experiment; and (iv) *Visual Elements* experiment, to capture user interactions with visualizations of varying types, complexity, dimensionality and appearance respectively. Furthermore, the tasks features different type using the taxonomy presented in Table 2. In particular, the Chart Type experiment consisted of simple tasks (e.g., comparison, retrieve value) across all visualization types using a single dimension. The Task Complexity experiment introduced more complex tasks such as finding data anomalies, computing derived values, or correlation. The Dimensionality experiment introduced tasks with 2, 3 and 4 dimensions. Tasks in the above experiments were delivered using default visual settings (i.e., no modifications on visual elements) and act as control tasks to the Visual Elements experiment. The tasks for the Visual Elements Experiment were divided into seven sets, each one introducing a visual element modification ( $ve_i \in VE$ ) as seen in Table 3.

#### 4.5 Study Procedure

Due to the implications of the national restrictions in response to the COVID-19 pandemic, it was decided that the study had to be conducted in a remote manner. Each participant was given access to the experiment platform, using the credentials used previously for extracting their user models.

The study was conducted for seven days where participants had to complete all four experiments (i.e., 160 visual exploration tasks). Once an experiment started, the platform loaded the analysis tasks in a random order and presented them to the participant in a consecutive fashion. A task was completed once the participant provided a response. One important constraint was that once an experiment started it could not be stopped until all tasks were addressed. During the experiment our system collected (i) the time (in milliseconds) taken by each participant to provide an answer to a specific task, and (ii) the validity of the answer provided.

At the beginning, all participants were given a set of training analysis tasks similar to those of the 4 experiments. Additionally, the users were given instructions on how to set their screen environment, such as setting the minimum screen size and screen resolution, to ensure that the study experience was the as identical as possible across different participants. This was essential to avoid problems such as avoiding cases where the visualization and/or task controls do not fit in a small screen causing the participant to scroll vertically, thus increasing their cognitive load and decreasing their performance. Moreover, prior to being able to engage with each of the actual experiments, written instructions were given to the participants indicating the overall experiment process and the approximate amount of time required to complete each experiment.

#### 4.6 Extracting Adaptation Rules

This section describes the process used for extracting adaptation rules for visualization types (*ARVT*) and visual elements (*ARVE*) from the recorded results of the study. While the collected results can provide insight on the influence of human factors when processing data visualizations it is not the goal of this work to go over such findings.

*Adaptation Rules for Data Visualization Types.* It must be noted that for generating the *ARVT* rule set, only responses from the three control experiments were used (i.e., responses for the Visual Elements experiment were excluded). The following five steps describe the procedure performed to extract rules for every human factor group. To facilitate our description, we will describe the procedure for a single group of participants, the ones with High Working Memory, as the process is identical for all other groups.

Step 1 starts by filtering the responses to the selected group of participants. In Step 2, the response time values are aggregated according to the task's metadata and chart type, to produce the average response time of each participant on every data visualization type, task types and data characteristics. Step 3, ranks the data visualization types for each participant at different task configuration levels, such as task type, and data characteristics. Ranking is based on the average response time the participant achieved when using a specific data visualization under a set of specific analysis task configuration. In Step 4 the results are further filtered to the data visualization with the highest Rank in combination with the analysis task configuration. The results are then summarized for all participants leading to six records, each representing the score of each of the six available data visualizations used for a specific analysis task configuration. Essentially, score in this context represents the number of times a specific data visualization had the best performance in terms of time response (in milliseconds) for a specific analysis task configuration across all participants. In the final step, the data visualization scores are normalized to the range [0..1], assuming that all rules represent homogeneous weak learners (i.e., all rules are of equal weight). The resulting *ARVT* is used by the adaptation engine to select the best fit data visualization as described in Section 3.

*Adaptation Rules for Visual Elements.* In this section, we define the rule extraction process of the *ARVE* set, which utilizes the responses captured from both the control experiments and the Visual Elements experiment. For instance, rule extraction for the "dark theme" visual element/setting requires the utilisation of responses from Visual Elements

521 experiment, which had "dark theme" enabled, and all corresponding analysis tasks from the rest of the experiments,  
522 acting as control analysis tasks. Similarly to the previous section, the *ARVE* rule set was generated in 5 steps. To  
523 facilitate the description, we consider the example of generating rules for the "dark theme" visual element high working  
524 memory participants on the six available data visualization types.

525  
526 Step 1 filters the responses so that only results from high working memory participants interacting with "dark  
527 theme" tasks are included. In Step 2, the response time variable is aggregated for each participant when dark theme is  
528 enabled and disabled across all data visualization types. Step 3 ranks the enable/disable options for each participant for  
529 each data visualization type based on the average response times. In Step 4, the results are further filtered to the data  
530 visualization with the highest Rank and are aggregated for all participants, resulting in the best dark theme option for a  
531 specific data visualization type. This results in 12 records (2 dark theme options  $\times$  6 visualization types) recording the  
532 number of times each option was the most effective for each visualization type. Finally, in step 5, the two dark theme  
533 options' scores are normalized to the range [0..1], similarly with the previous section. The resulting *ARVE* is used by  
534 the adaptation engine to configure the best fit data visualization elements as described in Section 3.  
535  
536  
537

## 538 5 ADAPTATION ENGINE EVALUATION STUDY

539  
540 In this section, we present a user study that evaluates the use of adapted/personalized data visualization to improve  
541 the efficiency and effectiveness of business data analysts. The study collects the following evaluation metrics: (i)  
542 performance and accuracy of participants when addressing visual analysis tasks; (ii) perceived user experience; and (iii)  
543 perceived system usability. For capturing the adaptation engine's impact on the above mentioned metrics we followed  
544 a within-subjects study design, enabling us to record all evaluation metrics when the participant navigated over:  
545 (i) the original non-adapted/personalized content, which includes analysis tasks with predefined data visualizations  
546 without any alterations or enhancements; and (ii) data analysis tasks that include dynamically adapted/personalized  
547 data visualizations generated from the adaptation engine. Our null hypotheses are the following:  
548  
549

550  
551 **Hypothesis H0<sub>1</sub>:** The performance of the participants in terms of milliseconds taken to address an analysis task,  
552 between the two conditions (i.e., adaptation enabled and disabled) will not be significantly different.  
553

554 **Hypothesis H0<sub>2</sub>:** The accuracy of the participants in terms of the total number of tasks addressed correctly, between  
555 the two conditions (i.e., adaptation enabled and disabled) will not be significantly different.  
556

557 **Hypothesis H0<sub>3</sub>:** The system's user experience score, between the two conditions (i.e., adaptation enabled and  
558 disabled) will not be significantly different.  
559

560 **Hypothesis H0<sub>4</sub>:** The system's usability score, between the two conditions (i.e., adaptation enabled and disabled) will  
561 not be significantly different.  
562  
563

### 564 5.1 Participants

565  
566 For this study we recruited 45 business data analyst participants of which 24 were male and 21 were female. Their age  
567 ranged between 25 and 60 years (mean age  $35.4 \pm 8.8$ ) and on average they had at least 2 years of experience in the  
568 field of data analytics. It has to be noted that 21 out of the 45 participants were also recruited in our previous user  
569 study described in Section 4.6. The new participants were recruited from the same organizations and adhered to the  
570 same participation preconditions as in the above mentioned user study. It must be noted that all new participants were  
571  
572

573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624

Table 7. Example pair of Non-adapted/Adapted Analysis Tasks

Task Name	Task Narrative	Task Type	Visualization Used
T01 Control Task	Identify the month with the highest sales during 2021 for brand "IdealCola".	Simple Comparison	Bar Chart
T01 Adaptive Task	Identify the month with the highest sales during 2021 for product "IdealCola Zero .33ltr x8 Can".		Adapted Data Visualization

given an appropriate time frame to go through our user modelling platform and perform all psychometric tests and questionnaires for the purpose of user modeling.

## 5.2 Dataset

The analysis dataset was constructed in collaboration with the partner organizations to cater for the diverse expertise and experience of the participant business analysts. A real dataset focusing on soft drinks sales was selected to be used for the study. The dataset comprises of 19 attributes representing the time, product and outlet dimensions (8 categorical, 1 boolean, 1 date, 4 integers and 5 measures) and consists of 731,446 real sale transactions recorded over the period of 3 years (2019-2021). The outlet names, brands and product names of the dataset were anonymized upon request of the providing organization.

## 5.3 Analysis Tasks

Using the realistic sales dataset, 38 analysis tasks were created in pairs of equal complexity, resulting in 19 pairs of tasks. An example of a task pair is presented in Table 7. Each pair includes: (i) the control non-adapted/non-personalised task addressed using a specific/predefined visualization type; and (ii) the adapted task which will utilise the adaptation engine’s capabilities to return the best fit data visualization to the participant addressing that particular task. Moreover, each pair of tasks has a specific analysis task type, which follows the same taxonomy of tasks presented in Table 2. In the background, the system records the time it took for the participant to interact and respond to the task question (in milliseconds) and assess whether the response of the participant was correct.

## 5.4 Study Material

Besides recording performance and accuracy for each task, the study aimed to capture the users’ experience and system usability factors. To this end, we utilised two accredited system evaluation questionnaires which we combined into a web-based questionnaire. Specifically for measuring the participants’ user experience we used the User Experience Questionnaire Sort Version (UEQ-S) [32]. According to the questionnaire’s authors, this questionnaire’s scales “cover a comprehensive impression of user experience. Both classical usability aspects (efficiency, perspicuity, dependability) and user experience aspects (originality, stimulation) are measured”. Moreover, for measuring the system’s usability we used the System Usability Scale (SUS) questionnaire [28]. This 10-scale questionnaire provides a reliable tool for measuring the usability of a system. We chose this tool as its deemed appropriate for our purpose since it has become an industry standard, with references in over 1300 articles and publications [28].

## 5.5 Study Procedure

Due to the implications of the national restrictions in response to the COVID-19 pandemic, it was decided that the study had to be conducted in a remote manner. At the beginning, all participants were invited into a remote MS Teams meeting where the team introduced the study and its goals. According to the dataset theme, it was decided that during the study participants would be assigned the role of a Brand Manager employed at a soft drinks company that sells products of the brand IdealCola.

A training session was conducted, allowing participants to experience the system, understand how its interface works and address demo analysis tasks. Towards the end of the training session the participants were given time to ask any questions that they had regarding the study. Once the training session was over a recorded version of the presentation and system demonstration videos was sent to all participants so they could revisit the training material covered. After the workshop, the team sent out emails to each individual participant including the URL to the platform, and personal user credentials to access it. This served as a second training step, so that each participant could login during their free time and practice an additional set of demonstration analysis tasks. This ensured that all participants were well familiarized to the platform's interface and analysis tools prior to engaging in the study. Moreover, it must be noted when a participant logged in for the first time was greeted with a welcome presentation of slides that demonstrated all the features of the data exploration system and also GDPR-related features (e.g., enabling a participant to request the deletion of their information). This presentation of features was also made available later during the actual study in case the participant wished to revisit it. Finally, since the study was conducted in a remote fashion, we needed to cater for external participant distractions. Appropriate controls were developed to prevent the system from tracking the user's view time/performance in the cases where the participant was distracted. In particular, if the participant was found to be non-responsive for 30 seconds the platform prompted the user to validate that they are still present, otherwise the current progress and all recorded statistics were reset and the user had to restart the task. The 30-second interval was decided with the industry partners after performing some tests with tasks of varying complexity.

The rest of the study was divided into two parts. The first part utilized the 19 non-adapted set of analysis tasks (i.e., the system returned the same predefined data visualization for the specific analysis task to all participants). The second part utilized the 19 adapted analysis tasks (i.e., the data visualizations were adapted by the adaptation engine, according to the user model and task characteristics). Finally, once a participant was finished with a part, they also had to complete the online questionnaires for assessing the system's user experience and usability factors. For each part, a participant could navigate in the list of available tasks (presented in random order), study their question/narrative and select one that will serve as the current task. The participant could then navigate to the Analysis Wizard, which is a visual analysis component enabling participants to perform data explorations. The functionality of the Analysis wizard was limited to providing the capabilities required to address the study tasks. During the exploration process, the participant was required to go through the three steps of the Analysis Wizard: (i) select analysis; (ii) select attributes; and (iii) view result (i.e., adapted or non-adapted visualization). Furthermore, the current analysis task's narrative was available to users on the top of the Analysis Wizard to remind them of what is required. Finally, once the participant had an answer for the analysis task, they could navigate back to the list of analysis tasks and provide their answer. During the exploration process, the platform was monitoring the interaction time of participants for all steps and assessed the correctness of their responses.

## 6 EVALUATION RESULTS

The following sections summarize the effectiveness of the delivered adaptation by reporting the impact of data visualization adaptation on the participant's performance, accuracy, perceived user experience and system usability factors as stated by the Hypotheses in Section 5.

### 6.1 Impact on Analysis Task Performance

The analysis of performance consisted of comparing the time required for participants to address paired (non-personalized vs. personalized) analysis tasks. In the beginning, the results were filtered to include only pairs where the participant responded accurately to both analysis tasks. An outlier analysis for assessing each participant's and each task's response times to reveal abnormal observations. Our analysis revealed for Task 13, which featured computation of a derived value, all participants took an extreme amount of time to complete the personalized task. After discussion with participants, it was understood that the majority of the participants were not aware on how to approach the analysis. Additionally, some participants reported having to use a calculator to find the correct answer. Consequently, it was decided that the task results were removed from the analysis.

The analysis across the two study conditions revealed that adaptation had a positive effect on participants' performance enabling them to achieve an average decrease of  $8.1 \pm 6.9$  seconds with regards to task completion time. Moreover, with adaptation enabled, performance improved for an average of  $9 \pm 2$  tasks per participant, while the number of tasks improved in terms of performance at the unique participant level was at maximum 15 tasks and at minimum 5 tasks. Additionally, with adaptation enabled, performance worsen for an average of  $2 \pm 1$  tasks per participant, while the number of tasks worsen in terms of performance at the unique participant level was at maximum 5 tasks and at minimum 0 tasks.

Analysis on the impact of adaptation with regards to performance across different task types shows that adaptation had a positive effect on participants' performance enabling them to achieve (i) a statistically significant average decrease of 7.8 seconds for Retrieve Value tasks ( $p < .01$ ), (ii) a statistically significant average decrease of 25.9 seconds for Correlation tasks ( $p = .01$ ), (iii) a statistically significant average decrease of 8.2 seconds on Simple Comparison tasks ( $p < .01$ ) and (iv) a non-statistically significant average decrease of 10.6 seconds on Compute Derived Value tasks ( $p = 0.24$ ). Since Simple Comparison tasks was the larger group of analysis tasks (10 task pairs), we decided to further explore this group of tasks by independently analysing Simple Comparison tasks which used time series data. Results show that with adaptation enabled, participants achieved (i) a statistically significant average decrease of 9.9 seconds on Simple Comparison tasks which used time series data ( $p < .01$ ) and (ii) a statistically significant average decrease of 4.5 seconds on the remaining Simple Comparison tasks ( $p < .01$ ). Moving on, with adaptation enabled performance improved for an average of  $84 \pm 82$  task responses across all analysis task types, while the number of task responses improved in terms of performance at the unique analysis task type level was at maximum 199 responses for Simple Comparison tasks which used time series data, and at minimum 5 responses for Compute Derived Value tasks. Finally, we report that with adaptation enabled, performance worsen for an average of  $22 \pm 21$  task responses across all analysis task types, while the number of task responses that worsen in terms of performance at the unique analysis task type level was at maximum 52 responses for Simple Comparison tasks which used time series data and at minimum 1 response for Correlation and Compute Derived Value tasks. Unfortunately, the sample of our tasks was limited to a single Find Anomaly pair of tasks, for which the majority of participants only responded correctly to the personalised variant of the task, leaving a very small sample of only four responses that could not be incorporated in this analysis.

In conclusion, the above results reveal that when the adaptation engine was enabled, participants' performance was positively affected and thus, we reject the null hypothesis  $H0_1$ .

## 6.2 Impact on Analysis Task Accuracy

The analysis of accuracy considered the participants' ability to address (non-personalized vs. personalized) analysis tasks correctly. For each of the study conditions a participant was able to achieve a maximum score of 19 (i.e., the total number of tasks).

Analysing the accuracy scores of each participant reveals that 62% of participants were more accurate when addressing analysis tasks with adapted/personalised data visualizations. Moreover, 18% of participants were not affected in terms of accuracy across the two study conditions, while the remaining 20% of participants were negatively impacted by adaptation in terms of accuracy. In contrast to analysis tasks with no data visualization adaptation, participants were able to address on average an additional 8% of analysis tasks correctly when working with tasks delivering data visualization adaptation. Analysis of accuracy scores across task types for both conditions revealed that participants were generally much more accurate in addressing tasks when adaptation was enabled for Simple Comparison, Compute Derived Value and Find Anomaly tasks. Specifically, participants were more accurate by 6.6% for Simple Comparison tasks, 34.2% for Computer Derived Value tasks and 90% for Find Anomaly tasks. In contrast, for Correlation and Retrieve Value task types we were not able to see a significant impact in terms of accuracy when participants were using adapted/personalised data visualizations for addressing the analysis tasks.

In conclusions, the overall analysis of task accuracy revealed that when the adaptation engine was enabled, participants' accuracy was positively affected and thus, we reject the null hypothesis  $H0_2$ .

## 6.3 Impact on User Experience and System Usability

During the evaluation study 35 participants responded to the User Experience Questionnaire (UEQ-S) [32] and the System Usability Score (SUS) questionnaire [28], right after they addressed all control analysis tasks (i.e., those tasks with a predefined/non-adapted data visualizations). Additionally, the same 35 participants responded to the two questionnaire right after they had address all analysis tasks for which the system produced an adapted data visualization. The collected data was analysed by an automated process offered by the questionnaire's authors in order to investigate  $H0_3$  and  $H0_4$ .

The User Experience Questionnaire has in total 8 scales, 4 scales measuring Pragmatic Quality (a metric that focuses on the task-oriented nature of an experience e.g., considers the task's efficiency and ease of use), and 4 scales measuring Hedonic Quality (a metric that focuses more on the fun, appeal and more generally on the originality aspects of the experience offered by a system). Using the responses of all participants we calculated Cronbach's alpha (or coefficient alpha) for each set of scales belonging to each metric (i.e., pragmatic quality and hedonic quality) for data collected in both conditions (i.e., adaptation disabled/enabled). Alpha values for both metrics across the two conditions were higher than 0.7 which is considered acceptable. Generally, scales that belong to the same group should show a high correlation. Therefore, using the Cronbach's alpha which is a measure for the consistence of a scale [8] helped us ensure that the different scales of the questionnaire were interpreted as intended by the participants. The baseline scores (i.e., adaptation disabled condition) for (i) pragmatic quality was 1.35, (ii) hedonic quality was 0.86 and (iii) the overall user experience was 1.11. Moreover, with adaptation enabled the score for (i) pragmatic quality was increased to 1.45, (ii) hedonic quality was increased to 0.97 and (iii) the overall user experience was increased to 1.21. The user experience scores achieved by the system across the two conditions are above the value of 0.8 and thus are considered a positive evaluation [32]. The adaptation engine evaluation revealed that enabling data visualization adaptation when

participants interact with the given analysis tasks facilitated an increase of their perceived user experience and thus we reject the null hypothesis **H0<sub>3</sub>**.

The baseline score (i.e., adaptation disabled condition) for System Usability was 66.2 with a standard deviation of 12.9. Moreover, with adaptation enabled the score for System Usability was increased to 67.4 with a standard deviation of 11. The usability scores achieved by the system across the two conditions (i.e., adaptation disabled/enabled) are considered marginally acceptable since a usability score is considered above average if it is higher than 68 [22]. Nonetheless, it is interesting to see that the usability score was increased by 1.2 when participants were exposed to the tasks utilising the adapted/personalised data visualizations. Another important aspect is that the standard deviation of the system usability score for when the adaptation was enabled is lower than the standard deviation of the baseline system usability score. A smaller standard deviation indicates that the average usability scores elicited from the participants' responses are closer to the mean and thus we are more confident for the higher system usability score achieved for the enabled adaptation condition. The adaptation engine evaluation revealed that enabling data visualization adaptation when participants interact with the given analysis tasks facilitated an increase of their perceived usability reaching to a marginal value closer to that of the average score and thus we can reject the null hypothesis **H0<sub>4</sub>**.

## 7 DISCUSSION

Our evaluation user study shows that the current adaptation engine and the adopted rule generation procedure improves the participants' performance and accuracy across a variety of data analysis tasks. Additionally, the platform positively affects the participants' perceived user experience and perceived system usability scores. Furthermore, the fuzzy rule-based classification framework used by the adaptation engine enables the quickly integrate new rules based on new data visualization interaction data. Additionally, the fuzzy adaptation logic and the ensemble processing approach used by the adaptation engine makes it easier to combine and utilise in parallel multiple adaptation driving factors (e.g., human factors) which might interact with each other. Moreover, the extraction of data visualization interaction data for adaptation rule generation, can take place during a user study using appropriate data collection tools, similar to the ones used in this work, or can be extracted from other sources, such as data analysis tools offering user interaction logging. Overall, the flexibility of the current framework allows for quickly collecting interaction data and efficiently transform it in adaptation rules. The framework is also open enough to be used for experimentation in other domains.

While our work committed to the improvement of the overall efficiency and effectiveness of the business data analyst when addressing data analysis tasks, there are some limitations that we would like to address in the future. The sample of analysis tasks used during evaluation was not balanced in terms of task type since more focus was given on simpler comparison tasks. Moreover, this work assumes homogeneous weak learners adaptation rules. An alternative approach would be to use a boosting framework by identifying which adaptations/interventions had the most influence on the improvement in terms of accuracy and performance. Some questions rising from this work that we plan in addressing as part of future endeavours includes: (i) How could our approach offer a transparent explanation to the business analyst as with regards to why the best-fit data visualization was selected? (ii) How can we more effectively process the resulting user's interaction with the adapted output and further gain insight on which adaptation/intervention was the most helpful for that type of user? and (iii) How does our adaptation perform with unexplored data visualizations and analysis task types? Our goal is to attempt to address these questions by first extending our sample of users and gathering more data visualization interaction data that can yield more diverse adaptation rules, thus, facilitating further exploration of the interaction of human factors on data visualizations, but also the exploration of this interaction as a driving force to the current adaptation engine.

## 8 RELATED WORK

Designing a user adaptive system involves the consideration of three questions; *what to adapt to*, *when to adapt* and *how to adapt* [3]. Our work turns the focus on *what to adapt to* and further explores *how to adapt* aspects. With regards to *what to adapt to*, research on information visualization reveals that individual differences influence how a user interacts, understands and utilizes data visualizations for performing analytical tasks. In fact, the growing interest on the effect of individual differences in information visualization resulted in comprehensive survey publications on the subject [19]. A subset of works on individual differences and their effect on information visualization includes and is not limited to the exploration of human factors such as cognitive abilities [4, 5, 17, 25, 29, 30, 33, 37], cognitive styles [16, 21, 26, 27, 31], personality traits [7, 14, 36, 37], and expertise/experience [17, 18, 29].

The question of *how to adapt* data visualizations is usually addressed (i) at the visualization type level i.e., using recommendations for a best fit data visualization, or (ii) at the individual visual element level i.e., applying modifications to, or additions of visual elements on a data visualization.

For instance, the works of Gotz et al. [12] and Grawemeyer [13] focused on adaptation with regards to delivering data visualization type recommendations based on user's interaction behavior or task features. On the other hand, visual element modifications are equally important to note. The work of Carenini et al. [3] investigated how the effectiveness of a data visualization (specifically a bar chart) can be increased with four different adaptive interventions. Additionally, in the context of Security Information and Event Management systems, Yelizarov et al. [35] proposed a graph of computer hosts that highlights the most significant hosts (i.e., graph nodes) and dims (using opacity) the rest according to the current cognitive load of the user for increasing efficiency when dealing with system threats. While Yelizarov et al. [35] leveraged the user's cognitive load for adaptation, others have utilised the underlying data for adapting the color of visual elements according to their mapped data category in order to reduce the user's cognitive load [24].

By contrast to the above works, our goal is to build a flexible human-centred by-design adaptation engine that leverages the power of a multidimensional human-centered user model for delivering the best fit data visualization (both in terms of data visualization type and visual element modifications). Specifically, our work targets data analyst users that perform visual data exploration in the context of a business environment, aiming to increase their comprehensibility of information leading to improved accuracy and time-to-action efficiency.

## 9 CONCLUSIONS AND FUTURE WORK

The adaptation engine evaluated in this paper adopts a fuzzy rule-based classification framework consisting of a fuzzy rule generation procedure and a classification procedure which selects the best fit data visualization for a given user performing a specific analysis task in two steps: (i) selecting the best data visualization type and (ii) the data visualization element modifications to be applied. Through this paper we presented the architecture of the adaption engine, and also the rule extraction process as this was performed using data visualization interaction data that was captured during a user study.

The evaluation of the adaptation engine using realistic data and 45 business data analysts revealed that the majority of participants were positively affected by the delivered data visualization adaptation in terms of their ability to correctly address analysis tasks. Additionally, we found that data visualization adaptation enabled our participants to execute their tasks significantly faster. This latter effect of performance improvement was more evident for simpler types of analysis tasks (e.g., Simple Comparison and Retrieve Value tasks). Moreover, the results show that the perceived user experience and system usability factors before and after adaptation was enabled were improved.

885 In the future, we plan to apply the proposed framework to other industry domains, generating a larger sample of  
886 more diverse adaptation rules which will further enable us to challenge our adaptation approach with more complex  
887 data visualizations and adaptive interventions. We also aim to investigate additional factors that contribute to the  
888 understandability and comprehension of data visualizations, such as transparency and explainability.  
889

## 890 ACKNOWLEDGMENTS

891 This research is partially funded by the Cyprus Research and Innovation Foundation under the projects IDEALVis  
892 (EXCELLENCE/0918/0366) and the European Union under the project SLICES-DS (No.951850).  
893  
894

## 895 REFERENCES

- 896 [1] Robert A. Amar, James R. Eagan, and John T. Stasko. 2005. Low-Level Components of Analytic Activity in Information Visualization. In *INFOVIS*.  
897 [2] Christos Amyrotos, Panayiotis Andreou, and Panagiotis Germanakos. 2021. *Human-Centred Persona Driven Personalization in Business Data*  
898 *Analytics*. Association for Computing Machinery, New York, NY, USA, 175–180. <https://doi.org/10.1145/3450614.3462241>  
899 [3] Giuseppe Carenini, Cristina Conati, Enamul Hoque, Ben Steichen, Dereck Tokar, and James Enns. 2014. Highlighting Interventions and User  
900 Differences: Informing Adaptive Information Visualization Support. In *Proc. ACM SIGCHI (Toronto, Ontario, Canada) (CHI '14)*. Association for  
901 Computing Machinery, New York, NY, USA, 1835–1844. <https://doi.org/10.1145/2556288.2557141>  
902 [4] Chaomei Chen. 2000. Individual differences in a spatial-semantic virtual environment. *Journal of the American Society*  
903 *for Information Science* 51, 6 (March 2000), 529–542. [https://doi.org/10.1002/\(SICI\)1097-4571\(2000\)51:6<529::AID-ASIS>3.0.CO;2-](https://doi.org/10.1002/(SICI)1097-4571(2000)51:6<529::AID-ASIS>3.0.CO;2-F)  
904 [F](https://doi.org/10.1002/(SICI)1097-4571(2000)51:6<529::AID-ASIS>3.0.CO;2-F) arXiv:[https://arxiv.org/abs/https://doi.org/10.1002/\(SICI\)1097-4571\(2000\)51:6<529::AID-ASIS>3.0.CO;2-F](https://arxiv.org/abs/https://doi.org/10.1002/(SICI)1097-4571(2000)51:6<529::AID-ASIS>3.0.CO;2-F)  
905 [5] Chaomei Chen and Mary Czerwinski. 1997. Spatial ability and visual navigation: an empirical study. *New Review of Hypermedia and Multimedia* 3, 1  
906 (July 1997), 67–89. <https://doi.org/10.1080/13614569708914684>  
907 [6] Claudia Imhoff and Colin White. 2011. Self-Service Business Intelligence Empowering Users to Generate Insights. [http://docs.media.bitpipe.com/io\\_](http://docs.media.bitpipe.com/io_10x/io_106625/item_583281/TDWI_Best_Practices_Report_Self-Service_BI_Q311%5B1%5D.pdf)  
908 [10x/io\\_106625/item\\_583281/TDWI\\_Best\\_Practices\\_Report\\_Self-Service\\_BI\\_Q311%5B1%5D.pdf](http://docs.media.bitpipe.com/io_10x/io_106625/item_583281/TDWI_Best_Practices_Report_Self-Service_BI_Q311%5B1%5D.pdf). Online; accessed 26 January 2021.  
909 [7] Cristina Conati and Heather Maclaren. 2008. Exploring the Role of Individual Differences in Information Visualization. In *Proc. AVI (Napoli, Italy)*  
910 *(AVI '08)*. Association for Computing Machinery, New York, NY, USA, 199–206. <https://doi.org/10.1145/1385569.1385602>  
911 [8] Lee J. Cronbach. 1951. Coefficient alpha and the internal structure of tests. *Psychometrika* 16, 3 (1951), 297–334.  
912 [9] Andreas Demetriou, George Spanoudis, and Michael Shayer. 2013. Developmental intelligence: From empirical to hidden constructs. *Intelligence* 41,  
913 5 (2013), 744–749. <https://doi.org/10.1016/j.intell.2013.07.014>  
914 [10] S B G Eysenck, H J Eysenck, and Paul Barrett. 1985. A revised version of the psychoticism scale. *Personality and Individual Differences* 6, 1 (1985),  
915 21–29. [https://doi.org/10.1016/0191-8869\(85\)90026-1](https://doi.org/10.1016/0191-8869(85)90026-1)  
916 [11] Mina Ghodrati, Lida Bavandian, Mostafa Morady Moghaddam, and Atena Attaran. 2014. On the relationship between problem-solving trait and the  
917 performance on C-test. *Theory and Practice in Language Studies* 4, 5 (feb 2014), 1093–1100. [https://link.gale.com/apps/doc/A373678951/AONE?u=](https://link.gale.com/apps/doc/A373678951/AONE?u=google scholar&sid=bookmark-AONE&xid=0e85444e)  
918 [google scholar&sid=bookmark-AONE&xid=0e85444e](https://link.gale.com/apps/doc/A373678951/AONE?u=google scholar&sid=bookmark-AONE&xid=0e85444e)  
919 [12] David Gotz and Zhen Wen. 2009. Behavior-Driven Visualization Recommendation. In *Proceedings of the 14th International Conference on Intelligent*  
920 *User Interfaces (Sanibel Island, Florida, USA) (IUI '09)*. Association for Computing Machinery, New York, NY, USA, 315–324. [https://doi.org/10.1145/](https://doi.org/10.1145/1502650.1502695)  
921 [1502650.1502695](https://doi.org/10.1145/1502650.1502695)  
922 [13] Beate Grawemeyer. 2006. Evaluation of ERST – An External Representation Selection Tutor. In *Diagrammatic Representation and Inference*, Dave  
923 Barker-Plummer, Richard Cox, and Nik Swoboda (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 154–167.  
924 [14] Tear M. Green and Brian Fisher. 2010. Towards the Personal Equation of Interaction: The impact of personality factors on visual analytics interface  
925 interaction. In *Proc. VAST (Salt Lake City, UT, USA)*. 203–210. <https://doi.org/10.1109/VAST.2010.5653587>  
926 [15] James J Gross and Oliver P John. 2003. Individual differences in two emotion regulation processes: Implications for affect, relationships, and  
927 well-being. , 348–362 pages. <https://doi.org/10.1037/0022-3514.85.2.348>  
928 [16] Marta Koć-Januchta, Tim Höffler, Gun-Brit Thoma, Helmut Prechtel, and Detlev Leutner. 2017. Visualizers versus verbalizers: Effects of cognitive  
929 style on learning with texts and pictures – An eye-tracking study. *Computers in Human Behavior* 68 (March 2017), 170–179. [https://doi.org/10.](https://doi.org/10.1016/j.chb.2016.11.028)  
930 [1016/j.chb.2016.11.028](https://doi.org/10.1016/j.chb.2016.11.028)  
931 [17] Sébastien Lallé, Cristina Conati, and Giuseppe Carenini. 2017. Impact of Individual Differences on User Experience with a Real-World Visualization  
932 Interface for Public Engagement. In *Proc. UMAP (Bratislava, Slovakia) (UMAP '17)*. Association for Computing Machinery, New York, NY, USA,  
933 369–370. <https://doi.org/10.1145/3079628.3079634>  
934 [18] S. Lee, S. Kim, Y. Hung, H. Lam, Y. Kang, and J. S. Yi. 2016. How do People Make Sense of Unfamiliar Visualizations?: A Grounded Model of  
935 Novice's Information Visualization Sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 499–508. <https://doi.org/10.1109/TVCG.2015.2467195>  
936

- 937 [19] Zhengliang Liu, R. Jordan Crouser, and Alvitta Ottley. 2020. Survey on Individual Differences in Visualization. *Computer Graphics Forum* 39, 3 (July  
938 2020), 693–712. <https://doi.org/10.1111/cgf.14033> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14033>
- 939 [20] Solomon Negash. 2004. Business Intelligence. *Communications of the Association for Information Systems* 13 (Jan. 2004), 177–195. <https://doi.org/10.17705/1CAIS.01315>
- 940 [21] Richard Riding and Graeme Douglas. 1993. The effect of cognitive style and mode of presentation on learning perfor-  
941 mance. *British Journal of Educational Psychology* 63, 2 (June 1993), 297–307. <https://doi.org/10.1111/j.2044-8279.1993.tb01059.x>  
942 arXiv:<https://bpspsychub.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2044-8279.1993.tb01059.x>
- 943 [22] Jeff Sauro. 2022. Measuring Usability with the System Usability Scale (SUS). <https://measuringu.com/sus/>. Online; accessed 20 March 2022.
- 944 [23] Susanne G Scott and Reginald A Bruce. 1995. Decision-Making Style: The Development and Assessment of a New Measure. *Educational and*  
945 *Psychological Measurement* 55, 5 (oct 1995), 818–831. <https://doi.org/10.1177/0013164495055005017>
- 946 [24] Vidya Setlur and Maureen C. Stone. 2016. A Linguistic Approach to Categorical Color Assignment for Data Visualization. *IEEE Transactions on*  
947 *Visualization and Computer Graphics* 22, 1 (2016), 698–707. <https://doi.org/10.1109/TVCG.2015.2467471>
- 948 [25] Ben Steichen, Giuseppe Carenini, and Cristina Conati. 2013. User-Adaptive Information Visualization: Using Eye Gaze Data to Infer Visualization  
949 Tasks and User Cognitive Abilities. In *Proc. IUI* (Santa Monica, California, USA) (*IUI '13*). Association for Computing Machinery, New York, NY,  
950 USA, 317–328. <https://doi.org/10.1145/2449396.2449439>
- 951 [26] Ben Steichen and Bo Fu. 2019. Towards Adaptive Information Visualization - A Study of Information Visualization Aids and the Role of User  
952 Cognitive Style. *Frontiers in Artificial Intelligence* 2 (2019), 22. <https://doi.org/10.3389/frai.2019.00022>
- 953 [27] Ben Steichen, Bo Fu, and Tho Nguyen. 2020. Inferring Cognitive Style from Eye Gaze Behavior During Information Visualization Usage. In *Proc.*  
954 *UMAP* (Genoa, Italy) (*UMAP '20*). Association for Computing Machinery, New York, NY, USA, 348–352. <https://doi.org/10.1145/3340631.3394881>
- 955 [28] SUS. 2022. System Usability Scale (SUS). <https://www.usability.gov/how-to-and-tools/methods/system-usability-scale.html>. Online; accessed 05  
956 March 2022.
- 957 [29] Dereck Toker, Cristina Conati, Giuseppe Carenini, and Mona Haraty. 2012. Towards Adaptive Information Visualization: On the Influence of User  
958 Characteristics. In *Proc. UMAP*, Judith Masthoff, Bamshad Mobasher, Michel C. Desmarais, and Roger Nkambou (Eds.). Springer Berlin Heidelberg,  
959 Berlin, Heidelberg, 274–285.
- 960 [30] Dereck Toker, Cristina Conati, Ben Steichen, and Giuseppe Carenini. 2013. Individual User Characteristics and Information Visualization: Connecting  
961 the Dots through Eye Tracking. In *Proc. ACM SIGCHI* (Paris, France) (*CHI '13*). Association for Computing Machinery, New York, NY, USA, 295–304.  
962 <https://doi.org/10.1145/2470654.2470696>
- 963 [31] N. Tsianos, P. Germanakos, Z. Lekkas, C. Mourlas, and G. Samaras. 2009. Eye-Tracking Users' Behavior in Relation to Cognitive Style within an  
964 E-learning Environment. In *Proc. ICALT* (Riga, Latvia). IEEE, 329–333. <https://doi.org/10.1109/ICALT.2009.110>
- 965 [32] UEQ-S. 2022. User Experience Questionnaire Sort Version (UEQ-S). <https://www.ueq-online.org>. Online; accessed 05 March 2022.
- 966 [33] M. C. Velez, D. Silver, and M. Tremaine. 2005. Understanding visualization through spatial ability differences. In *Proc. IEEE VIS* (Minneapolis, MN,  
967 USA). IEEE, 511–518. <https://doi.org/10.1109/VISUAL.2005.1532836>
- 968 [34] Herman A Witkin. 1971. *A Manual for the Embedded figures tests*. Consulting Psychologists Press, Palo Alto, Calif.
- 969 [35] Anatoly Yelizarov and Dennis Gamayunov. 2014. Adaptive Visualization Interface That Manages User's Cognitive Load Based on Interaction  
970 Characteristics. In *Proc. VINCI* (Sydney NSW, Australia) (*VINCI '14*). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/2636240.2636844>
- 971 [36] C. Ziemkiewicz, R. J. Crouser, A. R. Yauilla, S. L. Su, W. Ribarsky, and R. Chang. 2011. How locus of control influences compatibility with visualization  
972 style. In *Proc. IEEE VAST* (Providence, RI, USA). IEEE, 81–90. <https://doi.org/10.1109/VAST.2011.6102445>
- 973 [37] Caroline Ziemkiewicz and Robert Kosara. 2009. Preconceptions and Individual Differences in Understanding Visual Metaphors. *Computer Graphics*  
974 *Forum* 28, 3 (July 2009), 911–918. <https://doi.org/10.1111/j.1467-8659.2009.01442.x> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2009.01442.x>

975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988