

A methodology to automatically translate user requirements into visualizations: Experimental validation

Ana Lavalle^{a,b,*}, Alejandro Maté^{a,b}, Juan Trujillo^{a,b}, Miguel A. Teruel^{a,b}, Stefano Rizzi^c

^a Lucentia Research, Department of Software and Computing Systems, University of Alicante, Carretera San Vicente del Raspeig s/n, 03690, San Vicente del Raspeig, Alicante, Spain

^b Lucentia Lab, C/Pintor Pérez Gil, N-16, 03540, Alicante, Spain

^c DISI, University of Bologna, V.le Risorgimento 2, 40136, Bologna, Italy

ARTICLE INFO

Keywords:

Data visualization
Big data analytics
Model-driven development
Requirements engineering
Experimental validation

ABSTRACT

Context: Information visualization is paramount for the analysis of Big Data. The volume of data requiring interpretation is continuously growing. However, users are usually not experts in information visualization. Thus, defining the visualization that best suits a determined context is a very challenging task for them. Moreover, it is often the case that users do not have a clear idea of what objectives they are building the visualizations for. Consequently, it is possible that graphics are misinterpreted, making wrong decisions that lead to missed opportunities. One of the underlying problems in this process is the lack of methodologies and tools that non-expert users in visualizations can use to define their objectives and visualizations.

Objective: The main objectives of this paper are to (i) enable non-expert users in data visualization to communicate their analytical needs with little effort, (ii) generate the visualizations that best fit their requirements, and (iii) evaluate the impact of our proposal with reference to a case study, describing an experiment with 97 non-expert users in data visualization.

Methods: We propose a methodology that collects user requirements and semi-automatically creates suitable visualizations. Our proposal covers the whole process, from the definition of requirements to the implementation of visualizations. The methodology has been tested with several groups to measure its effectiveness and perceived usefulness.

Results: The experiments increase our confidence about the utility of our methodology. It significantly improves over the case when users face the same problem manually. Specifically: (i) users are allowed to cover more analytical questions, (ii) the visualizations produced are more effective, and (iii) the overall satisfaction of the users is larger.

Conclusion: By following our proposal, non-expert users will be able to more effectively express their analytical needs and obtain the set of visualizations that best suits their goals.

1. Introduction

Information visualization is paramount for the analysis of Big Data. The volume of data requiring interpretation is continuously growing. Visual analytics in software engineering is also gaining importance [1]. In fact, according to [2], the global data visualization market size stood at USD 8.85 billions in 2019 and is projected to reach USD 19.20 billions by 2027. The evolution of analytics and visualization techniques lies at the core of business strategies, and more and more research lines are focusing on the visualization of data.

However, users are typically unskilled in information visualization. Thus, finding the visualization that best suit a determined context is

a very challenging task for them. Moreover, it is often the case that users do not have a clear idea of what objectives they are building the visualizations for. Consequently, it is possible that graphics are misinterpreted, making wrong decisions that lead to missed opportunities. One of the underlying problems in this process is the lack of methodologies and tools that users who are not experts in visualizations can use to define their objectives and the corresponding visualizations.

Choosing and implementing the most suitable visualizations for each dataset is a really complicated task, particularly when working with Big Data. In these scenarios, it is common to find heterogeneous data sources that require extensive knowledge of the underlying data

* Corresponding author at: Lucentia Research, Department of Software and Computing Systems, University of Alicante, Carretera San Vicente del Raspeig s/n, 03690, San Vicente del Raspeig, Alicante, Spain.

E-mail address: alavalle@dlsi.ua.es (A. Lavalle).

<https://doi.org/10.1016/j.infsof.2021.106592>

Received 30 November 2020; Received in revised form 26 March 2021; Accepted 27 March 2021

Available online 5 April 2021

0950-5849/© 2021 The Authors.

Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

to create a suitable visualization [3]. Moreover, using an unsuitable type of visualization can lead to misunderstanding the data and making wrong decisions. In this sense, an approach such as SkyViz [4] can support users in creating visualizations. In SkyViz, the suitable visualization types for a given dataset are selected and created based on a *visualization context* defined by users; however, as the authors recognize, defining a visualization context from scratch can be a challenge for users who are not expert in data visualization.

To fill this gap, in this paper we present a process that helps non-expert users define their analytical goals and derive automatically the suitable visualizations according to the defined context. Our proposal covers the whole process, from the definition of the user requirements to the implementation of the visualizations. In our previous work we proposed (i) a User Requirements Model [5] to capture the users' analytical needs, (ii) a Data Profiling model [6] to extract semi-automatically the characteristics of the data sources, and (iii) a Data Visualization Model [6] that enables users to specify the visualization details regardless of the technology used for the implementation. Therefore, by following our proposal non-expert users will be able to communicate their analytical needs and obtain the visualizations that best suit them to achieve their goals. Besides, a dashboard will also be generated to group the visualizations and help users to carry out strategic decisions as such as the monitoring and measuring of their goals. Moreover, in this paper we put into practice the proposed methodology, by applying it to a case study focused on the Incidents Management from the Police Department of San Francisco.

To assess the validity of our proposal we have performed an experiment with 97 non-expert users in data visualization. In the experiment each user was tasked with two exercises. In the first exercise, participants were tasked with carrying out an analysis over a dataset without following any particular methodology. In the second exercise, each participant carried out a different analysis than the one they had seen before, this time following our proposed methodology. The results obtained from the experiment have been analyzed and represented graphically in order to show the improvements achieved by our methodology.

Therefore, the main contributions of this paper are to show the overall steps of the process, the application of the approach to a new scenario (illustrative example) to show its generalizability, and the validation of the proposal through the analysis of the results obtained by several groups of participants.

The rest of the paper is structured as follows. Section 2 presents the related work in the area of visualizations and analytics. Section 3 describes our process to automatically create visualizations. Section 4 shows our approach applied to an illustrative example. Section 5 presents an evaluation of the proposal by means of an experiment with non-expert users. Section 6 describes the validity threats to our proposal. Finally, Section 7 summarizes the conclusions and sketches future works.

2. Related work

Several approaches highlight the importance of visual analytics. For instance, [7] and [8] show the potential of visual analytics in software engineering. In [7] a visualization framework is presented that utilizes heat-maps to explore the evolution of a source code repository. Meanwhile, [8] presents visualization approach that captures significant aspects of the development process, and then tightly integrates and synchronizes them with product artifacts created by it.

Due to the relevance of this field, numerous authors are working in this area. In [9,10] and [11], techniques are proposed to automatically generate visualizations or dashboards. However, all of them rely on the user to choose the type of visualization to be used. That is why some other approaches propose ways to find the best type of visualization. For instance, authors in [12] review the main classifications proposed in the literature and integrate them into a single framework. In [13]

a framework is proposed that chooses the best type of visualization. Similarly, in [14] some visualization types are related to those types of users objectives that could be more compliant with. Finally, the SkyViz approach asks users to specify a structured visualization context and determines the suitable types of visualization [4].

Other works are focused on the possible limitations of graphic representations. [15] argues that one of the reasons for the lack of advanced visualizations are users, who often do not know how they may represent their data. Similarly, in [16], the authors point out that users are often seen as the “weakest link” in the security chain. For this reason, the authors propose an approach that improves systems by ensuring that problems are mitigated even when the users deviate from their expected behavior. In [17] a classification of causes of pitfalls is proposed, where pitfalls are responsibility of either the designer or the user. They list three types of (negative) effects: *cognitive*, *emotional*, and *social*. The distinction between designer and user-induced mistakes is particularly valuable in pragmatic terms, as it can give immediate insights to the producers or to the evaluators of visualizations respectively. In this sense, visualization designers should look at the encoding of the visualization, while users should pay attention to pitfalls in the decoding.

It is crucial to consider the possible risks and errors that can be made during the design and generation of visualizations. [18] points out that the rendering process introduces uncertainty in three areas: *data collection process*, *algorithmic errors*, and *computational accuracy and precision*. Moreover, in [19] the authors presented an initial study about the representation of errors and uncertainties visually. The possible sources of uncertainty are acquisition, model, transformation and visualization.

It is also relevant for users to understand the visualization that they are seeing and what is the goal that this visualization pursues. Visualizations are required to be precise and easy to comprehend by users in order to minimize the interpretation errors made by users as well as designers. Visualizations must also contemplate the changing needs of users, considering high-level semantics, and reasoning about unstructured and structured data, providing easier access and better understanding of the data [20]. Moreover, although often overlooked in visualization design, requirement modeling is a paramount activity [21] that compensates for the little attention usually paid to (explicitly) representing the reasons, i.e., the *why*, in terms of motivations, rationale, objectives, and requirements.

Despite all the work done in this field, none of the approaches previously mentioned provides a methodology that guides non-expert users from the start in the specification of the most adequate set of visualizations and facilitates their generation and grouping into suitable dashboards used for the extraction of knowledge. In this sense, our proposal aims to better bridge the gap between user requirements and visual analytics.

3. Process to create visualizations automatically

In this section we describe our methodology. Fig. 1 represents the proposed process. The first model is the **User Requirements Model** presented in [5]. The main aim of this model is to capture the users' analytical needs. Since we are dealing with non-expert users, the model is completed by following a sequence of guidelines. Then, a **Data Profiling Model** [5] is obtained. This model is created by semi-automatically analyzing the features of the data sources selected in the previous model. Once both models are completed, a **Visualization Specification** is derived according to their information. This specification contains enough information to automatically derive the suitable visualization types by following [4]. This transformation generates the **Data Visualization Model** [6], which allows users to specify visualization details regardless of the underlying technology used for the implementation. Using this model, users are also able to confirm whether the proposed visualizations fulfill the essential requirements for which they were created and whether they contribute to reach the users goals by

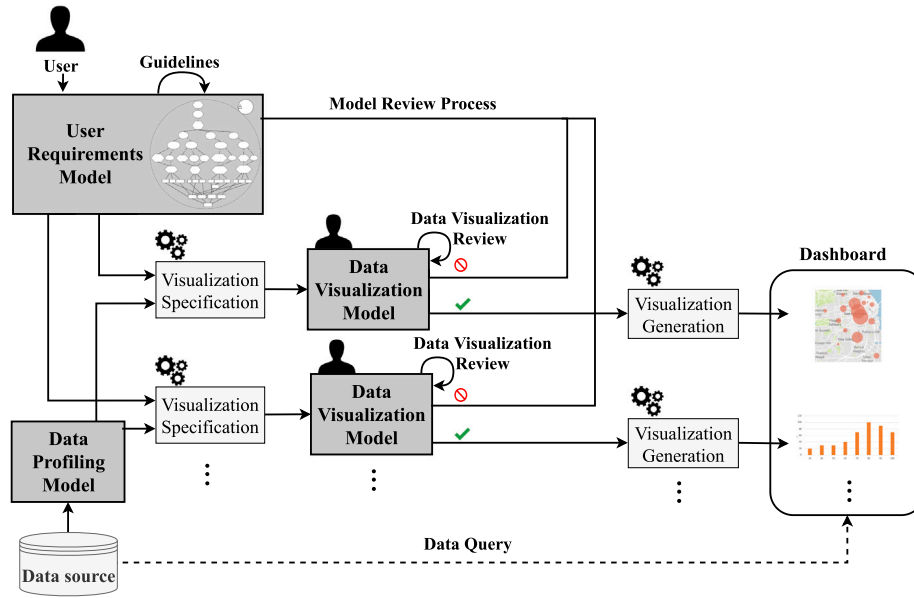


Fig. 1. Process proposed to create visualizations.

providing the necessary answers or not. If a visualization does not pass the validation, it means there are missing or ill-defined requirements. In this case, the models will be reviewed to identify which aspects were not taken into account. Otherwise, if a visualization passes the validation, it will be implemented in the selected technology and added to a dashboard.

In the following sections we explain in more detail the different components of the process.

3.1. User Requirements Model

The User Requirements Model supports the users in the definition of their data analysis objectives and helps to determine which visualization types they need to achieve these objectives. This model collects the **Interaction** and **Visualization Goals** that compose the Visualization Specification. Section 4.1 shows this model applied in an illustrative example.

In order to formally define our model, in [5] we proposed the metamodel shown in Fig. 2. This metamodel is an extension of [22], used for social and business intelligence modeling, and derived from the widely known *i**, in its 2.0 version [23] and its specialized *i** for Data Warehouses extension [24]. *i** is one of the most widespread frameworks and has been successfully applied to a large number of fields, such as [25–27]. Moreover, it facilitates the communication with the user, structures the information (objectives and mechanisms to achieve them) in an intuitive way, and provides a structure to the requirements.

Elements from *i** are represented in blue, elements from *i** for Data Warehouses in red, and the new concepts added in yellow. In the following we describe in detail the main elements of the metamodel.

- **Visualization Actor:** the user who will interact with the system. It can be classified as either *Tech* or *Lay* depending on whether she is expert or not in complex data visualizations.
- **Business Processes:** the process at the core of users' analysis. It serves as a guideline for the definition of their *Goals*.
- **Strategic Goals:** the main objectives of the business process; achieving them translates into an improvement from a current situation into a better one.
- **Analysis Type:** it allows users to express which kind of analysis they wish to perform, as classified by [28]:

- **Prescriptive:** How to act?
- **Diagnostic:** Why has this happened?
- **Predictive:** What is going to happen?
- **Descriptive:** What to do to make it happen?

- **Decision Goals:** decisions aimed at taking appropriate actions to fulfill a strategic goal. They also explain how the associated strategic goal can be achieved.
- **Information Goals:** the lower-level abstraction goals that represent the analysis to be carried out over the available information.
- **Visualization:** a specific visualization type that will be implemented to satisfy one or more information goals.
- **Visualization Goals:** they describe the data aspects that the visualization tries to reflect. Work in [5] proposed a flowchart to aid users in finding which visualization goal they are pursuing. The flowchart contains a series of Yes/No questions to be answered by users, and provides an easy way to discern which visualization goals should be included for each visualization. The possible goals that users can choose from are [4]:

- **Composition:** Highlight how the parts of data are composed to form a total.
- **Order:** Order values.
- **Relationship:** Analyze correlation.
- **Comparison:** Establish similarities and dissimilarities.
- **Cluster:** Emphasize the grouping into categories.
- **Distribution:** Analyze how data are dispersed in the space.
- **Trend:** Examine the general tendency.
- **Geospatial:** Analyze data using a geographical map.

- **Interaction Type:** Type of interaction that the visualization must support. In [5] a series of guidelines was proposed to help users choose one or more types of interaction they want to be supported by the visualization. The possible interactions that users can choose from are [4]:

- **Overview:** Gain an overview of the entire data collection.
- **Zoom:** Focus on items of interest.
- **Filter:** Quickly focus on interesting items by eliminating unwanted items.
- **Details-on-demand:** Select an item and get its details.

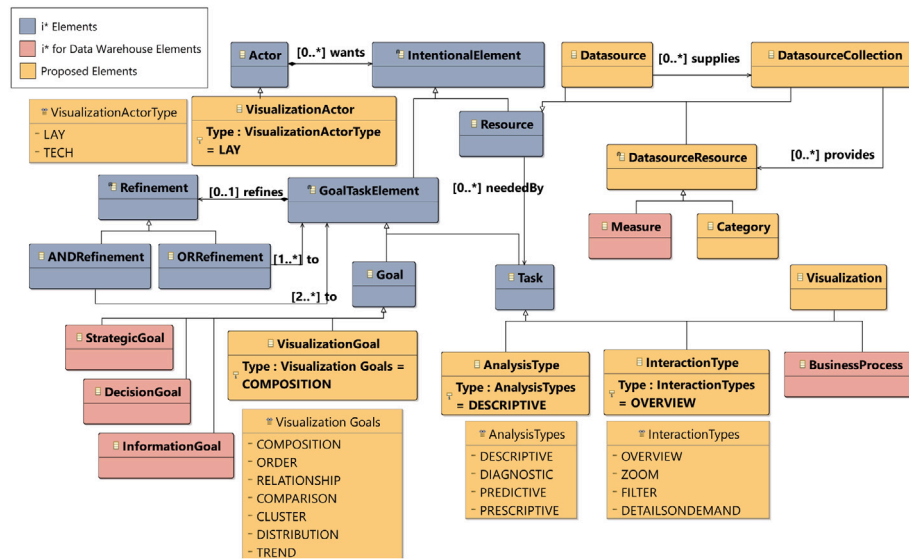


Fig. 2. Visualization specification metamodel.

- **Datasource Resource:** elements that provide relevant data from the data source.

3.2. Data Profiling Model

Following the proposed process, the next model is the Data Profiling Model. At first, in the User Requirements Model, the users have captured the data elements to be represented in each visualization. Then, the Data Profiling Model captures the data characteristics that are relevant for that visualization, such as **Dimensionality**, **Cardinality**, and **Dependent/Independent Type**. In [5], a Java implementation of a Data Analyzer to carry out data profiling was described. This software allows users to specify the data source from which they need to extract information and performs the extraction in an automated and guided way. Section 4.2 shows an example.

These characteristics are extracted in a semi-automatic manner, as explained below. First, the users specify a connection to the source dataset they wish to visualize. Then, a menu is provided where users can choose if they wish to retrieve the Data type, Cardinality, or Dimensionality of the selected column. Finally, the software returns the information requested by users. This tool has been created to collect information about the data in a simple way for users. In order to know how to delimit the values for each coordinate we have followed the approach proposed in [4], which classifies the Dimensionality, Cardinality, and Dependent/Independent Type as follows:

- **Dimensionality** is used to declare the number of variables to be visualized. Specifically, it can be *1-dimensional* when the data to represent is a single numerical value or string, *2-dimensional* when one variable depends on another, *n-dimensional* when a data object is a point in an n-dimensional space, *Tree* when each item in a collection is linked to one parent item, or *Graph* when each item in a collection is linked to an arbitrary number of items.
- **Cardinality** represents the number of data items. It is set to *Low* if this number is below a few dozens, to *High* otherwise.
- **Type of Data** is used to declare the type of each variable v . We identify each category as follows. If v is numerical, it is labeled as *Interval* if it supports the determination of equality of intervals or differences, as *Ratio* if it also has a unique and non-arbitrary zero point. If v is alphanumeric, the program shows a list of values; the user can then specify if in the list there is an order (in which case v is *Ordinal*) or not (*Nominal*).

3.3. Visualization Specification

Once the User Requirements Model and the Data Profiling Model are completed, the information coming from the models composes the Visualization Specification. We follow the SkyViz approach to discover which type of visualization suits bests each particular case, taking into account users preferences. Section 4.3 shows an example.

As described in [4], SkyViz operates by (i) asking the user to define a visualization context based on seven prioritizable coordinates for assessing her objectives and describing the dataset to be visualized; (ii) translating the visualization context into a set of suitable visualization types; (iii) asking the user to select one preferred visualization type among those proposed at the previous step; (iv) finding the best bindings between the columns of the dataset and the graphic coordinates used by the visualization type chosen by the user, and (v) asking the user to select one preferred binding among those proposed at the previous step. Specifically, as to (i), the seven coordinates composing the visualization context are filled starting the User Requirements Model and the Data Profiling Model. Step (ii) is performed based on a *suitability function* that assesses to which extent (fit, acceptable, discouraged, unfit) each visualization type is suitable for each possible value of each visualization coordinate; for instance, pie charts are discouraged for high-cardinality data, and bubble graphs are fit for n-dimensional data. The scores in the suitability function were mainly derived from the best practices found in the literature [29–31]. The set of suitable visualization types is then defined as those that are Pareto-optimal; a visualization type is Pareto-optimal when no other visualization type dominates it, being better along one coordinate and not worse along all the other coordinates. Given one preferred visualization type among the Pareto-optimal ones, step (iv) requires to decide how each variable in the dataset will be visualized, i.e., to establish a binding between each variable and each graphic coordinate of that visualization type. This is done by relying on a set of scores that indicates to which extent each graphic coordinate of each visualization type is suitable for each data type; for instance, the ‘X’ graphic coordinate of a single line chart is fit for variables of interval and ratio type, and the ‘size’ graphic coordinate of a bubble graph is unfit for variables of nominal type. Like for step (ii), the bindings proposed to the user are all the Pareto-optimal ones.

In [6] we explain in detail how to transform the Visualization Specification into a visualization following a Model-Driven Architecture (MDA) standard. As Fig. 3 shows, we transform the Visualization Specification by means of a set of model-to-model transformations using the QVT language [32], a standard from the OMG. For example, to

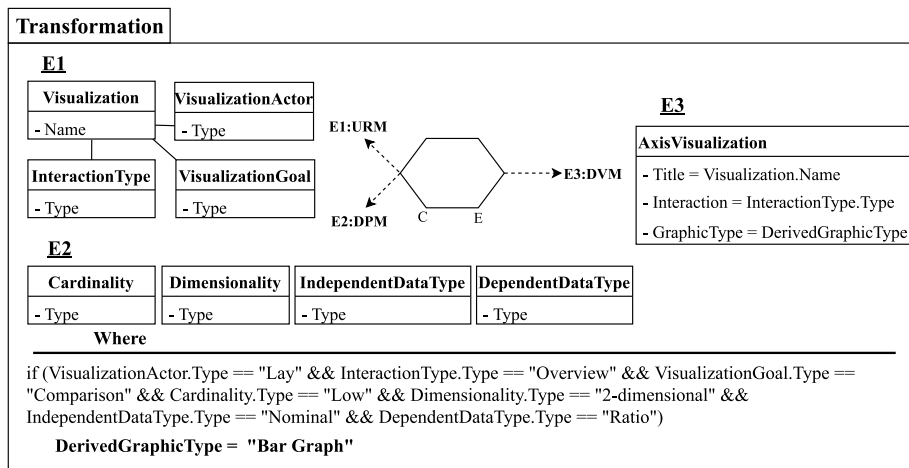


Fig. 3. Transformation of Visualization Specification into a visualization type.

derive an axis-based visualizations, our transformation generates an AxisVisualization element according to the graphic type established by the transformation. To derive this value we use the imperative part of the transformation (Where clause) according to the specific criteria established by [4] for each graphic type. The *Cardinality*, *Dimensionality*, *IndependentDataType*, and *DependentDataType* values are obtained from the **Data Profiling Model**, while the *VisualizationActor*, *InteractionType*, and *VisualizationGoal* are obtained from the **User Requirements Model**.

3.4. Data Visualization Model

Once a suitable visualization is obtained, we allow users to customize it using the Data Visualization Model [6]. This model shows a representation of the visualization and facilitates non-expert users in selecting the right visual analytics. For instance, users may define the element that determines the order in the visualization, the orientation of the visualization, or other elements as the legend, font family, and range of colors. Section 4.4 shows an example.

Users will modify the visualization through the Data Visualization Model until it meets their requirements. Once all the elements have been customized, users have to confirm whether the visualization obtained achieve their goals. If the visualization passes the validation, it will be generated as described in the next subsection. Otherwise, an unsuccessful validation generates a review of the existing User Requirements Model. The review process consists in reviewing the User Requirements Model to add or modify missing goals. If the visualization does not meet the goal for which it was created, the users will have the opportunity to redefine that goal or create a new one.

3.5. Visualization Generation

If the visualization achieves the requirements for which it was created, it will pass the validation and will be generated. This step transforms each element specified in the Data Visualization Model into a code-level specification for a graphic library. The transformation is done by means of code-to-text transformations that generate the code according to the target library. In our case, we use either the D3 JavaScript [33] or the Plotly [34] libraries to render the visualizations. The visualization, combined with the other ones generated in the process, will be grouped into a dashboard so that the user has access to all the information simultaneously. Section 4.5 shows an example of the generation of a visualization.

4. Illustrative example

This section shows the approach applied to an illustrative example based on the Police Department Incident Reports dataset [35] from the open dataset of San Francisco city [36]. In this case, the Police Department of the city requires a set of visualizations to analyze their data in order to help them improve the responsiveness of their services and reduce the incidents. We have applied our proposal to this case study by following the process in Fig. 1.

4.1. User Requirements Model

The first element involved in our process is the User Requirements Model (Section 3.1); Fig. 4 shows the result of its application. In this case, the final user is the Police Department Supervisor of the city of San Francisco, represented as a “Lay user” because she is not a specialist in visualization of Data Analytics. The Business Process which the user wants to analyze is “Incidents Management”, and the strategic goal that she wishes to achieve is to “Reduce incidents”.

In order to achieve this strategic goal, the user decides to perform a “Prescriptive analysis” and decomposes it into two decision goals, “Identify risk of the incident” and “Identify workload of police districts”, that aim to fulfill the strategic goal.

Afterwards, the user specifies information goals for each decision goal. These goals represent the lowest level of goal abstraction. In the case of decision goal “Identify risk of the incident”, the user refines it into two information goals, “Analyze neighborhoods with more incidents” and “Analyze the categories of incidents”. Decision goal “Identify workload of police districts” is refined it into the information goal “Analyze the number of incidents attended by police district”.

At this moment, the user has the essential information about her goals, and she can start to define the visualization context. For each information goal, a visualization will be automatically derived in order to achieve it. Each visualization represents one or more visualization goals (aspects of the data the visualization is trying to reflect) and one or more kinds of interaction (how users would like to interact with the visualization). A set of guidelines that may be used by users to aid in the definition of these elements can be found in [5]. In this case the user has selected for the different visualizations “Distribution”, “Geospatial”, “Comparison”, and “Trend” as visualization goals and “Overview” as interaction type.

Finally, visualizations are decomposed into Categories and Measures that will populate them. In this case, the visualization of “Number of incidents by neighborhood” includes “Neighborhoods” as category, and “Amount incidents” as measure. For the visualization of “Number of incidents by category” the user picked “Incident category” as category,

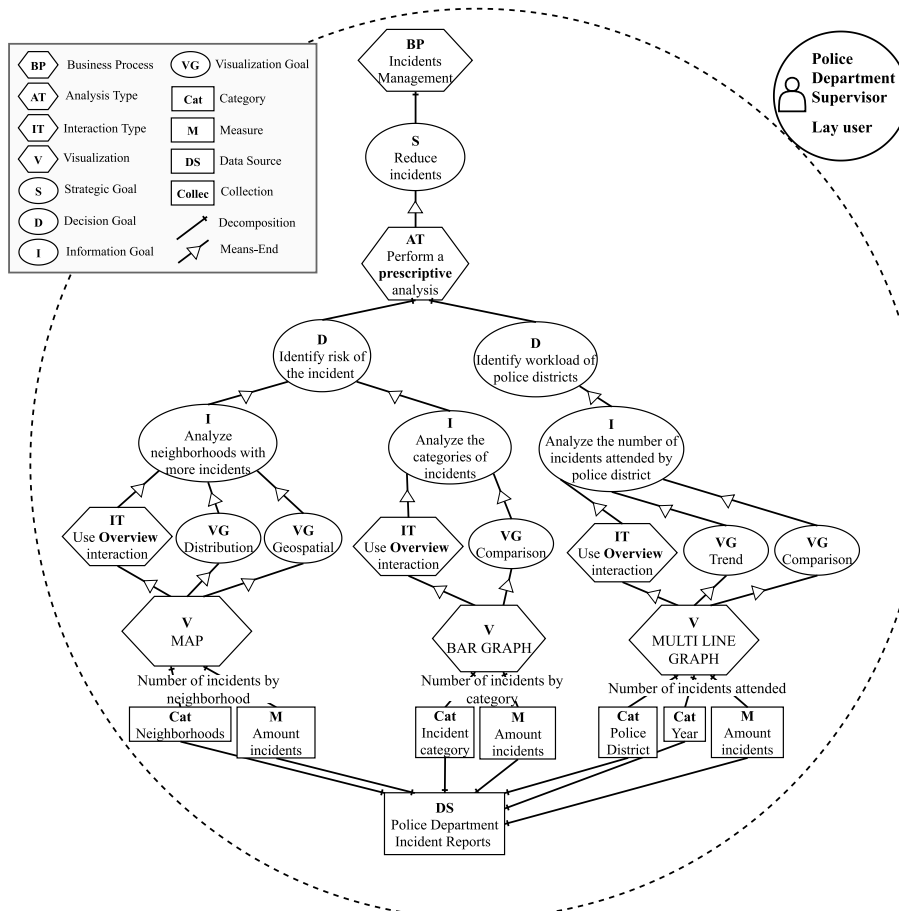


Fig. 4. Application of our User Requirements Model to the case study.

and “Amount incidents” as measure. Finally, in the case of “Number of incidents attended”, it makes use of the categories “Police District” and “Year” and the measure “Amount incidents”.

4.2. Data Profiling Model

Once the data sources and collections that will feed the visualizations have been defined by the user, we apply the Data Profiling Model (Section 3.2) this model determines, in a semi-automatic way, the Dimensionality, Cardinality, and Dependent/Independent Type of the data. We focus on the “Number of incidents by category” visualization from the User Requirements Model, which requires information about category “Incident category” and measure “Amount incidents”.

First, through the Data Profiling Model, the independent variable “Incident category” is classified as *Nominal* and the dependent variable “Amount incidents” as *Ratio*. Dimensionality is set as *2-dimensional*, because the user has selected 2 variables to visualize. Finally, the Cardinality is defined as *Low* because the independent variable contain 19 items to represent.

4.3. Visualization Specification

Once the Visualization Specification has all the necessary information from the previous models, it is used as input of the approach presented in [4]. This approach performs a suitability function that assesses to which extent (fit, acceptable, discouraged, unfit) each visualization type is suitable for the information stored in the Visualization Specification (Section 3.3).

Table 1 shows the Visualization Specification with its suitability scores (though for brevity we only include three visualization types, all

Table 1
Suitability scores for different visualizations types.

Vis. specification	Bar graph	Bubble graph	Single line graph
Goal: Comparison	<i>fit</i>	<i>fit</i>	<i>unfit</i>
Interaction: Overview	<i>fit</i>	<i>fit</i>	<i>fit</i>
User: Lay	<i>fit</i>	<i>acceptable</i>	<i>fit</i>
Dimensionality: 2-dim.	<i>fit</i>	<i>unfit</i>	<i>fit</i>
Cardinality: Low	<i>fit</i>	<i>acceptable</i>	<i>acceptable</i>
Independent type: Nominal	<i>fit</i>	<i>unfit</i>	<i>unfit</i>
Dependent type: Ratio	<i>fit</i>	<i>fit</i>	<i>fit</i>

the available visualization types were actually compared). According to the suitability scores, the most suitable visualization for the case at hand is “Bar Graph”. Fig. 3 shows how we use transformations to automate this process.

4.4. Data Visualization Model

Once the visualization type has been established as “Bar Graph”, a Data Visualization Model (Section 3.4) is built as Fig. 5 shows to verify that the visualization satisfies the users’ needs and allow them to customize it.

This model shows a mockup of the visualization with a series of characteristics that the user can customize. For example, the user has selected “Amount incidents” for the X axis, “Incident category” for the Y axis, and the orientation has been determined as *Horizontal*. When the user has finished customizing the visualization, she will have to test if the visualization makes it possible to satisfy the information goal “Analyze the categories of incidents” (i.e., if all the necessary information

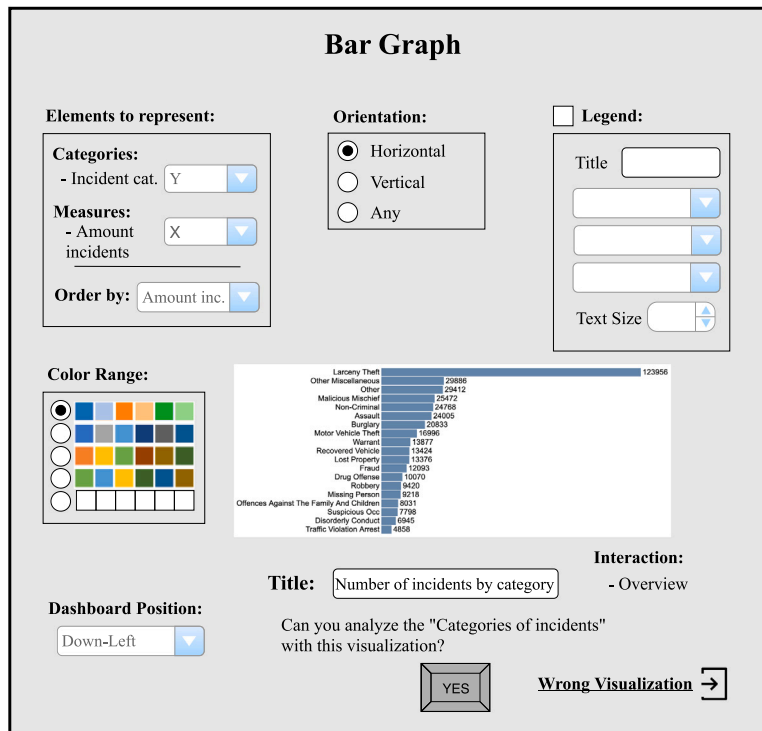


Fig. 5. Data Visualization Model.

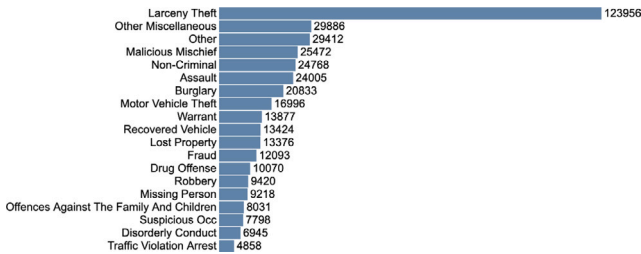


Fig. 6. Generated visualization.

can be analyzed). If the visualization passes the validation, it will be generated.

4.5. Visualization Generation

After the validation is passed, visual requirements are translated into an implementation by means of calls to the D3 JavaScript library [33] (Section 3.5), obtaining the visualization shown in Fig. 6. Consequently, this visualization, combined with those generated from the other informational goals “Analyze neighborhoods with more incidents” and “Analyze the number of incidents attended by police district” will be added to the dashboard that will enable non-expert users in data visualizations – such as the Police Department Supervisor – to monitor their processes.

5. Evaluation

In this section, we present the performance of our proposal in a controlled experiment. This experiment is part of a set of experiments for assessing the validity and impact of the proposal. In [37] it is possible to find a copy of the experimental materials in order to reproduce the experiments. We have followed the guidelines for experimentation in software engineering proposed in [38]. We have

carried out our experiments with non-expert users in data visualization coming from the University of Castilla la Mancha (UCLM) Campus of Albacete (Spain) and from a small IT company located in Alicante (Spain).

5.1. Experiment context

The main goal of these experiments is to analyze the proposed methodology and evaluate its understandability and effectiveness from the viewpoint of non-expert users in data visualization. In the experiments, a total of 97 non-expert participants filled in the questionnaires. The set of participants included 2nd-year computer engineering students and employees of a technological company. In both cases none of the participants had knowledgeable skill in data visualization.

The students were recruited through an email from their teachers, and participated voluntarily. They were rewarded with 0.25 out of 10 in the final mark of the subject, however, their performance had no impact on the mark. The participants of the company participated on a voluntary basis without any benefit.

Due to the COVID measures, not all the participants could meet in the same room and this is why they had to be divided into the groups shown in Table 2. The group of instructors was composed by two developers of the method and two professors from the University of Castilla la Mancha. The professors were instructed to know the experiment and what kind of assistance they could provide. In the case of the experiment in the company, the instructors were two developers of the method. During the experiment it was not explicitly explained who were the developers of the method.

As to the assistance provided during the experiments, it was focused on the development of the exercises, not on the content. The different elements of the model were explained so that users were able to generate the goal tree by themselves. Some additional help was provided to derive the visualizations, since the experiment was made on paper and the prototype CASE tool [39] was not ready at that time. Moreover, using the prototype would have introduced additional risks and noise, since it would have been difficult to understand whether an

Table 2
Participants to the experiment.

	No. of participants	Background	Assistance
Group 1	9	UCLM	Yes
Group 2	15	UCLM	Yes
Group 3	21	UCLM	Yes
Group 4	39	UCLM	No
Group 5	13	IT company	Yes

Table 3
Main features of the experiment.

Null hypothesis	- H_{0A} : The use of the proposed methodology does not allow users to cover more analytical questions - H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations - H_{0C} : Users do not find any improvement between performing the exercises with or without the proposed
Dependent variable	- Number of questions answered - Perceived value of the visualizations - Perceived improvement
Independent variable	- Whether the methodology was used or not - Whether there has been assistance in the performance of the experiment
Location	- Albacete - Alicante
Date	- October 2020
Subjects	- 84 Computer Engineering Students - 13 Employees of an IT company

improvement in the results was derived from the methodology itself or from the usage of the tool. In the case of the participants who had no assistance, there was no interaction at all between them and the instructors.

Importantly, only the help that the tool would have provided was indeed given to the participants. There was no help in applying the methodology, as this would have posed a threat to the validity of the experiment.

As [Table 3](#) shows, the experiment seeks to discover whether (i) the proposed methodology really helps in answering more analytical questions, (ii) it increases the perceived value of the set of visualizations created, and (iii) whether users perceive an improvement when doing an exercise with or without the methodology. Then, the independent variables were defined as (i) whether the methodology was used or not and (ii) whether there has been assistance to carry out the experiment or not. And finally, the null hypothesis that the experiment tried to accept/refuse.

5.2. Experiment design

The experiment consisted of performing two exercises related to a tax collection topic and an evaluation, the first exercise without following any methodology and the second by following our methodology. Usually different cases are used, however, in this experiment we decided to use the same case to avoid fatigue effect risk, since the experiment was very long.

Firstly, before starting the exercises, we requested the participants to fill a short anonymous survey where they were asked about their age, gender, studies, and level of experience with data visualization tools. In this way we could then identify non-expert users and evaluate how our proposal improves their results. Both the survey and the experiment exercises were always filled in an anonymous manner, making us unable to identify the author behind the survey and the corresponding exercise.

Then, users performed the requested exercises. On each exercise, participants were assigned with a different strategic goal to achieve related to the tax collection topic. In the first exercise, participants were

Table 4
 2×2 factorial design.

	Exercise 1 Without methodology	Exercise 2 With methodology
Experiment mode A	Strategic goal 1	Strategic goal 2
Experiment mode B	Strategic goal 2	Strategic goal 1

asked to define visualizations by knowing the strategic goal and having all the data available. In this first case, the participants did not follow any method. In the second exercise, participants were assigned with a different strategic goal and, in this case, they were asked to follow our methodology. Once both exercises were finished, the participants completed the evaluation by answering concrete questions that required the usage of the visualizations they had created. They also had to rate the visualizations they had defined as well as the improvement perceived when doing an exercise with or without the methodology.

Everyone did the experiment first without the method and then with it. To avoid the learning effect, a 2×2 factorial design with confounded interaction [40] was used, as shown in [Table 4](#). In this sense, the strategic goal to achieve and the analytical questions were swapped, i.e., the participants with Experiment Mode A received strategic goal 1 to do exercise 1 (without using our methodology) and strategic goal 2 to do exercise 2 (using our methodology). Conversely, in Experiment Mode B the participants received strategic goal 2 to do exercise 1 (without using our methodology) and strategic goal 1 to do exercise 2 (using our methodology). The experiment modes were distributed equally among the participants.

The analytical questions that participants had to answer by using the created visualizations are listed below. These questions were established by the authors during a brainstorming process. Moreover, the questions were tested in a pilot experiment and some of them were removed. For each question, participants must state whether they can or cannot answer it using their previously defined visualizations. It was not possible to answer all the questions with a single visualization since the questions were designed to force participants to use more than one visualization.

Reduce unpaid bills (Strategic Goal 1)

1. Identify the areas with most unpaid bills
2. Identify the types of taxes with most unpaid bills
3. Identify the tax records with most unpaid bills
4. Analyze the evolution over time of unpaid bills

Reduce the bill collection time (Strategic Goal 2)

1. Identify the amount of bills paid on and after the deadline
2. Identify the types of bills that are mostly paid after the deadline
3. Indicate in which areas there are payment delays
4. Identify the most delayed tax records

Finally, in order to rate their confidence on the visualizations created, they were tasked to fill the rubric shown in [Table 5](#). This table allows participants to communicate the perceived value of the set of visualizations created in Exercise 1 and Exercise 2 and the improvement perceived between performing the exercises with or without our methodology. This is a subjective aspect that allows us to know if users can really feel that there is an improvement in the performance of the exercise by following our methodology.

Therefore, the information collected was: (i) information regarding participants demographics, (ii) number of analytical questions answered, (iii) score of the rubric ([Table 5](#)), and (iv) time required by the participants to complete the experiment, which was only collected for statistical purposes.

Table 5
Rubric to evaluate the set of visualizations.

		Score			
		1	2	3	4
Ex.1	Are the visualizations useful? Is the information well represented? Are the visualizations suitable for the information?	Strongly disagree	Disagree	Agree	Totally agree
Ex.2	Are the visualizations useful? Is the information well represented? Are the visualizations suitable for the information?	Strongly disagree	Disagree	Agree	Totally agree
Ex.1 vs. Ex.2	Did you perceive any improvement in Ex. 2 over Ex. 1?	No improvement	Little improvement	Reasonable improvement	Remarkable improvement




Fig. 7. Histogram for number of questions answered for Group 123.

Table 6
Comparison of the analytical questions and rubric results based on whether or not the methodology was used.

		Average	
		Without methodology	With methodology
Group 123	Answered questions	1.87	2.72
	Perceived value	1.82	2.74
	Perceived improvement	2.95	
Group 4	Answered questions	2.13	2.34
	Perceived value	1.89	2.58
	Perceived improvement	2.42	
Group 5	Answered questions	1.08	2.15
	Perceived value	1.85	2.62
	Perceived improvement	2.92	




Fig. 8. Perceived value of visualizations for Group 123.

5.3. Experiment results

After manually transcribing the survey, the analytical questions, and the rubric for a subsequent analysis, we obtained the results shown in Table 6. We have grouped the results in:

- **Group 123:** 45 Computer Engineering students from UCLM who were assisted while carrying out the experiment.
- **Group 4:** 39 Computer Engineering students from UCLM who were not given assistance in carrying out the experiment.
- **Group 5:** 13 Employees of a small IT company who were assisted while carrying out the experiment.

In the following, the results from each group will be analyzed.

5.3.1. Group 123 - Students with assistance

Group 123 included 45 Computer Engineering Students from the University of Castilla la Mancha. In this case, we gave assistance to the participants through a detailed explanation of the methodology and solved all their doubts.

According to the results obtained, shown in Table 6, the set of visualizations generated without following any methodology can answer 1.87/4 (47%) of the specific questions proposed, while this number grows until 2.72/4 (68%) coverage when following the proposed method. Furthermore, a 2-Sample T-Test was performed with an alpha

of 0.05. Thanks to this test, we could conclude that the mean number of questions answered differs at the 0.05 level of significance, with a p -value < 0.001. Therefore, for Group 123, with a 95% confidence level, we can reject the null hypothesis “ H_{0A} - The use of the proposed methodology does not allow users to cover more analytical questions”. Thus, the number of analytical questions answered by using the methodology is significantly different (higher) from the number of analytical questions answered without using the methodology. Moreover, Fig. 7(a) reflects the normality of the data since it corresponds to the structure of a Gaussian distribution [41].

In order to accept or reject the null hypothesis “ H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations”, the answers of the rubric of Table 5 have been analyzed. As Table 6 shows, the participants scored the visualizations created without methodology with an average 1.82/4. Comparatively, the visualizations generated using the methodology presented were scored with an average of 2.74/4. Fig. 8 represents the perceived value of the resulting visualizations.

The 2-Sample T-Test was performed with an alpha of 0.05. Thanks to this test, we could conclude that in the case of the perceived value of the visualizations the means differ at the 0.05 level of significance, with




Fig. 9. Histogram for the confidence on the visualizations score for Group 123.

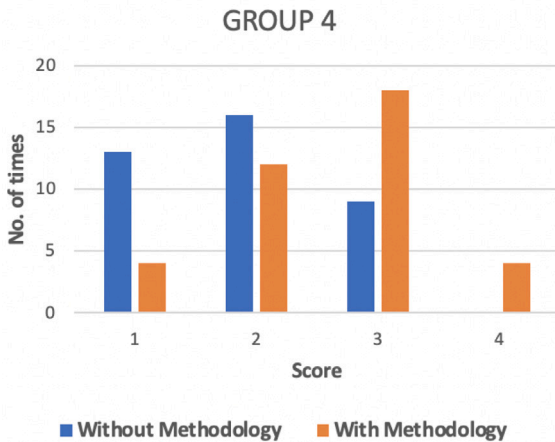


Fig. 10. Perceived value of visualizations for Group 4.

a p -value < 0.001 . Therefore, for Group 123, with a 95% confidence level, we can reject the null hypothesis “ H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations”, meaning that, for this group, the perceived value of the visualizations is indeed higher when using our methodology. As in the previous case, Fig. 9 reflects the normality of the data, as well as the difference of the averages. Therefore, the normality of our data is confirmed. We conclude that the results show a statistical significance that confirms the impact of the methodology proposed.

5.3.2. Group 4 - Students without assistance

Group 4 is composed of 39 Computer Engineering Students from the University of Castilla la Mancha. In this case we let them carry out the experiment without offering them any assistance.

In accordance with the results obtained (Table 6), the set of visualizations generated without following any methodology can answer 2.13/4 (53%) of the specific questions proposed, while this number grows until 2.34/4 (59%) when following the proposed method. However, in this case, the 2-Sample T-Test concludes that with a p -value of 0.430, the number of questions answered is not significantly different. Therefore, for “Group 4”, we cannot reject the null hypothesis “ H_{0A} - The use of the proposed methodology does not allow users to cover more analytical questions”.

In order to accept or reject the null hypothesis “ H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations”, the answers of the rubric of Table 5 have been analyzed. As Table 6 shows, the participants scored the visualizations created without methodology with an average 1.89/4. Comparatively, the visualizations generated using the methodology presented were scored with an average of 2.58/4. Fig. 10 represents the perceived value of the resulting visualizations.

The 2-Sample T-Test was performed with an alpha of 0.05. Thanks to this test, we can conclude that in the case of the perceived value of the visualizations the means differ at the 0.05 level of significance, with a p -value < 0.001 . Therefore, for Group 4, with a 95% confidence level, we can reject the null hypothesis “ H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations”, meaning

that in this group the perceived value of the visualizations is not the same using or not using our proposed methodology.

Fig. 11 reflects the normality of the data, as well as the difference of the averages. Therefore, the normality of our data is confirmed.

5.3.3. Group 5 - Company employees with assistance

The last group, number 5, was composed of 13 employees from the small technological company. In this case, we gave assistance to the participants through a detailed explanation of the methodology and solved all their doubts.

According to the results obtained, shown in Table 6, the set of visualizations generated without following any methodology can answer the 1.08/4 (27%) of the specific questions proposed, while this number grows until 2.15/4 (54%) coverage when following the proposed method. Furthermore, a 2-Sample T-Test was performed with an alpha of 0.05. Thanks to this test, we could conclude that in the case of the number of questions answered means differ at the < 0.05 level of significance, with a p -value of 0.019. Therefore, for “Group 5”, with a 95% confidence level, we can reject the null hypothesis “ H_{0A} - The use of the proposed methodology does not allow users to cover more analytical questions”, meaning that the number of analytical questions answered by using the methodology is significantly different (again, higher) from the number of analytical questions answered without using the methodology.

Fig. 12 reflects the normality of the data, as well as the difference of the averages. Therefore, the normality of our data is confirmed.

In order to accept or reject the null hypothesis “ H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations”, the answers of the rubric of Table 5 have been analyzed. As Table 6 shows, participants scored the visualizations created without methodology with an average 1.85/4. Comparatively, the visualizations generated using the methodology presented were scored with an average of 2.62/4. Fig. 13 represents the perceived value of the resulting visualizations.

The 2-Sample T-Test was performed with an alpha of 0.05. Thanks to this test, we could conclude that in the case of the perceived value of the visualizations the means differ at the 0.05 level of significance, with a p -value of 0.047. Therefore, for “Group 5”, with a 95% confidence level, we can reject the null hypothesis “ H_{0B} : The use of the proposed methodology does not improve the set of generated visualizations”, meaning that in Group 5 the perceived value of the visualizations is higher when following our proposed methodology.

Finally, as in the previous cases, Fig. 14 reflects the normality of the data, as well as the difference of the averages. Therefore, the normality of our data is confirmed.

In conclusion, the T-test results showed statistical significance for the results obtained, confirming the impact of the methodology proposed. Fig. 15 summarizes the score given to the improvement of one method over the other through the third question of rubric shown in Table 5. In most cases the participants have detected an improvement when using our methodology.

5.3.4. Analysis of visualizations

Finally, we analyzed the visualizations generated freely and those generated using our methodology.

The first outcome is that, by following our methodology, a larger number of visualizations were created than by creating them freely.




Fig. 11. Histogram for the confidence on the visualization score for Group 4.




Fig. 12. Distribution of data of number of questions answered for Group 5.

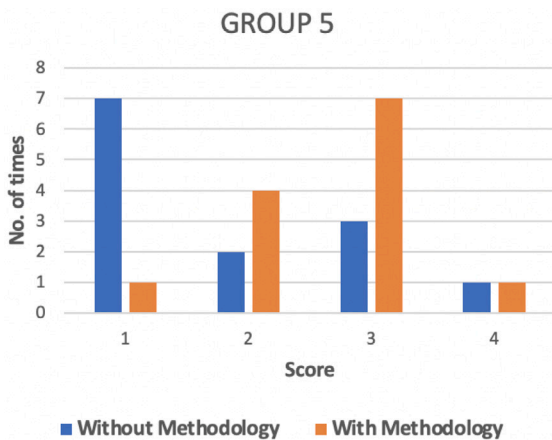


Fig. 13. Perceived value of visualizations for Group 5.

A total of 205 visualizations were created freely (an average of 2.11 per participant), while 257 visualizations were created by following our methodology (an average of 2.65 per participant).

Moreover, we have analyzed the types of visualizations selected in each case. When the participants did the exercise freely, the most used visualization types were Column Graph, Pie Chart, and Map as Fig. 16(a) shows. Nevertheless, when the participants did the exercise following our methodology, the most used visualizations were Column Graph, Map, and Bubble Graph as Fig. 16(b) shows.

Therefore, we can conclude that: (i) following the methodology, participants tend to use more visualizations; (ii) the visualization type most used by the participants is also the one most recommended by our methodology; and (iii) when participants use visualizations that are not suitable for non-expert users, such as histograms, it is common to create erroneous visualizations that do not really represent what they expected.

5.4. Meta-analysis

In this section the results from the different groups are discussed. Table 7 summarizes the results obtained.

The results increase our confidence about the utility of our methodology, because (i) it allows users to cover more analytical questions, (ii) it improves the set of generated visualizations, and (iii) users find improvements when they use it to execute the exercises. For the group that carried out the experiment without any assistance (Group 4), it was not possible to verify statistically that the use of the proposed methodology allows users to cover more analytical questions. Therefore, given the positive results obtained in the remaining groups, we can infer that

Table 7

Summarized results.

	Group 123	Group 4	Group 5
H _{0A} : The use of the proposed methodology does not allow users to cover more analytical questions	Rejected	Not rejected	Rejected
H _{0B} : The use of the proposed methodology does not improve the set of generated visualizations	Rejected	Rejected	Rejected
H _{0C} : Users do not find any improvement between performing the exercises with or without the proposed methodology	Rejected	Rejected	Rejected

some assistance or prior training is required for the effective application of the methodology. In addition, these results point in the direction of emphasizing the development and usage of a user-friendly tool to apply our proposal more effectively, reducing the users' knowledge burden and improving the results obtained.

6. Validity threats

In this section, we summarize the main limitations and validity threats for the performed experiments. Although we did our best to avoid that the outcome is affected by undesired factors, there are some aspects that must be taken into account when reproducing these experiments:

- When performing the experiments, we had a data analyst supporting non-expert users in order to aid in following the methodology. Such actor may not be always available, which may alter the results (i.e., Group 4). We are working to a user-friendly CASE tool to verify that users are able to define visualization requirements completely on their own.
- Our proposal is meant to be context-independent. We have applied it in educational, economic, smart cities, and gas turbine contexts. However, we have not applied our proposal yet in a full set of contexts, so there may be some specific user profiles we have not considered yet.
- Our methodology increases the capability to answer analytical questions. However, it is still recommended that the user who defines the visualizations is an expert in the application domain for which the visualizations are required.
- We rely on [4] to derive suitable visualization types. This means that our proposal inherits the associated limitations when deriving the visualizations. One of such limitations is that not all visualization types are supported. Furthermore, if a significantly




Fig. 14. Distribution of data of confidence on the visualization score for Group 5.




Fig. 15. Perceived improvement results.

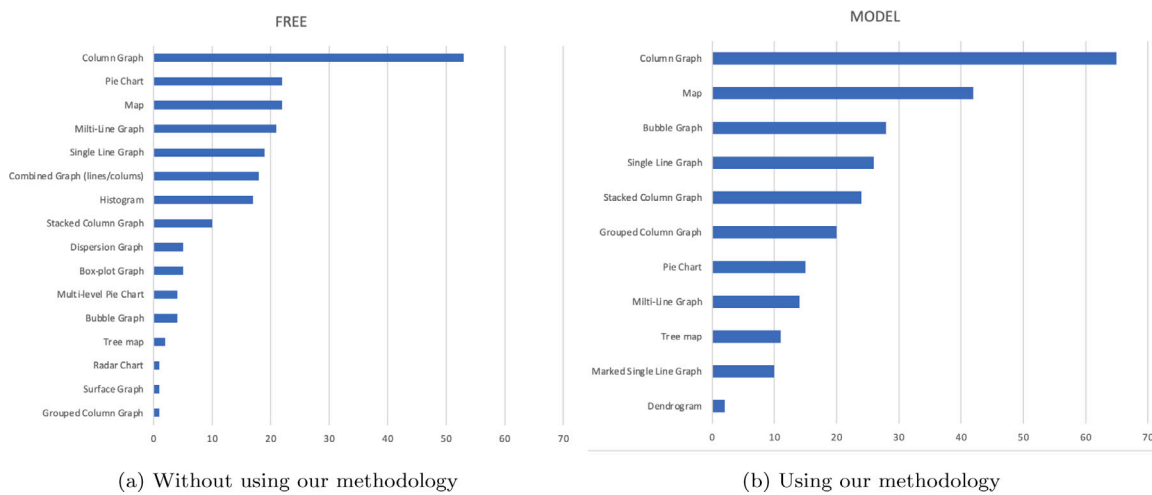


Fig. 16. Most used visualizations.

larger number of visualization types were to be included, the seven coordinates we rely on might no longer be sufficient to distinguish them.

- The participants in the experiment received a predefined template (essentially, a tree with empty nodes) as a guide to facilitate the creation of the User Requirements Model. Then they completed the model independently by filling the nodes and adding or eliminating branches as necessary.
- Although the objective of the experiment was to test our approach on non-expert users only, the experience of the users can be considered a validity threat.

7. Conclusions and future work

The volume of data that needs be analyzed and interpreted is continuously growing. Data visualization plays a key role in this analysis. However, finding the most effective visualizations is a difficult task. Normally, users are not experts in data visualization, and they rarely know which is the visualization type that will best suit them, nor they know exactly what information they are trying to extract

from them. Unfortunately, there is a lack of methodologies that guide non-expert users, taking into account their analysis goals to define the visualizations they need. For this reason, in this paper we have presented a process that helps non-expert users define their analytical goals and achieve them by automatically deriving the visualizations that best suit a certain context.

Compared to other approaches, our proposal covers the whole process, from defining user requirements to implementing visualizations. Therefore, the great advantage of our proposal is that non-expert users will be guided to reflect their analytical needs and automatically obtain a set of visualizations that will help them to achieve their goals.

To evaluate the impact of our proposal, we have presented a case study and performed a set of experiments with non-expert users in data visualization. The experiments have been carried out by 97 participants, including 84 Computer Engineering Students and 13 employees of a technological company, all of them non-expert in data visualization. These experiments confirmed the validity of our proposal since it has been shown that our methodology (i) allows users to cover more analytical questions, (ii) improves the set of generated visualizations and, (iii) users themselves perceive improvements when adopting our

Table 8
Experiment data 1.

Group	Id	Mode	No. of questions answered without methodology	No. of questions answered with methodology	Value of visualizations without methodology	Value of visualizations with methodology	Comparative	Time
1	1	A	0	4	1	4	4	49
1	2	A	2	3	2	4	4	49
1	3	A	2	4	3	3	4	56
1	4	A	1	4	2	3	3	54
1	5	B	1	3	1	3	3	53
1	6	B	2	2	1	2	2	
1	7	B	2	2	1	2	2	
1	8	B	3	4	2	4	4	49
1	9	B	1	4	1	3	4	47
2	1	A	4	4	2	3	4	54
2	2	A	1	4	3	4	3	58
2	3	A	2	2	1	1	2	61
2	4	A	1	4	2	4	4	60
2	5	A	1	3	1	3	3	60
2	6	A	1	4	2	3	3	68
2	7	A	3	3	3	2	2	63
2	8	A						62
2	9	B	2	4	1	4	4	63
2	10	B	2	2	3	4	2	62
2	11	B	4	3	3	3	2	62
2	12	B	1	1	3	3	3	59
2	13	B	2	1	1	3	3	59

Table 9
Experiment data 2.

Group	Id	Mode	No. of questions answered without methodology	No. of questions answered with methodology	Value of visualizations without methodology	Value of visualizations with methodology	Comparative	Time
2	14	B						57
2	15	B	3	1	2	3	4	54
3	1	A	2	4	1	2	2	44
3	2	A	1	2	2	3	2	45
3	3	A	3	3	2	1	2	47
3	4	A	1	0	3	3	2	62
3	5	A	3	2	2	1	2	60
3	6	A	1	3	2	3	4	50
3	7	A	3	2	3	3	3	50
3	8	A						50
3	9	A	2	0	2	1	1	51
3	10	A	2	3	3	4	2	51
3	11	A						55
3	12	B	3	3	1	2	4	62
3	13	B	2	3	2	3	3	57
3	14	B						57
3	15	B	0	2	1	2	4	55
3	16	B	3	3	1	3	3	53
3	17	B	2	2	2	2	3	50
3	18	B	0	2	1	1	2	49
3	19	B						47
3	20	B	2	4	1	2	4	47

methodology. Although the majority of user groups in the experiments have shown a statistical significance in favoring the methodology, for the group that carried out the experiment without any assistance it has not been possible to verify statistically that the use of the proposed methodology allows users to cover more analytical questions. The other improvements have also been confirmed with this group.

Therefore, considering that the assistance in following the method has a positive impact on its application, we are implementing a user-friendly tool [39]. As part of our future work we are going to test the usability of the tool through new controlled experiments. This will allow us to adjust the tool to users' needs. Moreover, we will explore the possibility of taking into account changing needs to our methodology.

CRediT authorship contribution statement

Ana Lavallo: Conceptualization, Methodology, Investigation, Data curation, Writing - original draft, Visualization. **Alejandro Maté:** Conceptualization, Methodology, Writing - review & editing, Supervision,

Project administration. **Juan Trujillo:** Conceptualization, Methodology, Resources, Writing - review & editing, Supervision, Project administration, Funding acquisition. **Miguel A. Teruel:** Formal analysis, Investigation, Data curation, Writing - review & editing. **Stefano Rizzi:** Conceptualization, Methodology, Writing - review & editing.

Declaration of competing interest

No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to <https://doi.org/10.1016/j.infsof.2021.106592>.

Acknowledgments

This work has been co-funded by the ECLIPSE-UA (RTI2018-0942-83-B-C32) project funded by Spanish Ministry of Science, Innovation. Ana Lavallo holds an Industrial PhD Grant (I-PI 03-18) co-funded by the University of Alicante, Spain and the Lucentia Lab Spin-off Company.

Table 10
Experiment data 3.

Group	Id	Mode	No. of questions answered without methodology	No. of questions answered with methodology	Value of visualizations without methodology	Value of visualizations with methodology	Comparative	Time
3	21	B	2	2	1	3	3	45
4	1	A	2	4	2	4	4	46
4	2	A	2	2	2	2	2	46
4	3	A	2	2	2	3	3	46
4	4	A	3	2	2	1	1	47
4	5	A	3	3	2	3	3	48
4	6	A	3	0	2	3	3	48
4	7	A	1	1	2	2	3	47
4	8	A	3	2	2	3	3	50
4	9	A	1	2	1	1	2	50
4	10	A	3	4	3	3	2	53
4	11	A	2	3	3	4	2	52
4	12	A	0	3	1	2	2	54
4	13	A	2	3	3	2	2	56
4	14	A	3	4	3	3	2	55
4	15	A	2	4	2	3	4	58
4	16	A	2	2	2	2	2	57
4	17	A	2	3	1	3	3	66
4	18	A	1	4	2	3	3	60
4	19	B	3	2	2	3	3	65
4	20	B	3	4	2	2	3	64
4	21	B	1	1	1	2	2	64

Table 11
Experiment data 4.

Group	Id	Mode	No. of questions answered without methodology	No. of questions answered with methodology	Value of visualizations without methodology	Value of visualizations with methodology	Comparative	Time
4	22	B	1	3	3	3	3	62
4	23	B						62
4	24	B	0	2	1	2	2	61
4	25	B	4	3	1	3	4	58
4	26	B	1	1	2	2	2	56
4	27	B	4	3	3	3	3	56
4	28	B	3	2	1	3	3	55
4	29	B	2	2	3	2	1	55
4	30	B	0	0	1	1	1	55
4	31	B	0	2	1	3	2	54
4	32	B	1	2	1	2	2	51
4	33	B	3	3	3	4	3	50
4	34	B	3	2	1	3	2	50
4	35	B	3	1	3	3	2	50
4	36	B	4	4	2	4	3	49
4	37	B	1	1	1	1	1	49
4	38	B	3	1	2	3	2	47
4	39	B	4	2	1	2	2	46
5	1	B	0	3	1	3	4	
5	2	B	1	3	1	3	3	
5	3	B	2	1	3	3	2	
5	4	B	1	2	2	3	2	

Table 12
Experiment data 5.

Group	Id	Mode	No. of questions answered without methodology	No. of questions answered with methodology	Value of visualizations without methodology	Value of visualizations with methodology	Comparative	Time
5	5	B	0	3	3	3	3	
5	6	B	1	2	1	2	3	
5	7	A	4	2	4	1	2	
5	8	A	1	3	2	2	3	
5	9	A	1	1	1	2	2	
5	10	A	1	0	1	2	4	
5	11	A	0	1	1	3	3	
5	12	A	1	3	3	3	3	
5	13	A	1	4	1	4	4	60

We would like to thank Elena Navarro, Pascual González and Víctor López from the University of Castilla-La Mancha (Spain) for their collaboration in the experiment.

Appendix. Experiment data

See [Tables 8–12](#).

References

- [1] M. Staron, H.A. Sahraoui, A. Telea, Special section on visual analytics in software engineering, *Inf. Softw. Technol.* 98 (2018) 117, <http://dx.doi.org/10.1016/j.infsof.2018.03.001>, URL <https://doi.org/10.1016/j.infsof.2018.03.001>.
- [2] R. Arterburn, Data visualization market growth, future prospects and competitive analysis (2020-2027) — fortune business insights, 2020, URL <https://fresnobserver.com/data-visualization-market-growth-future-prospects-and-competitive-analysis-2020-2027-fortune-business-insights/>.
- [3] E.G. Caldarola, A.M. Rinaldi, Improving the visualization of wordnet large lexical database through semantic tag clouds, in: *International Congress on Big Data, IEEE*, 2016, pp. 34–41.
- [4] M. Golfarelli, S. Rizzi, A model-driven approach to automate data visualization in big data analytics, *Inform. Vis.* 19 (1) (2020) <http://dx.doi.org/10.1177/1473871619858933>, URL <https://doi.org/10.1177/1473871619858933>.
- [5] A. Lavalle, A. Maté, J. Trujillo, S. Rizzi, Visualization requirements for business intelligence analytics: A goal-based, iterative framework, in: *27th IEEE International Requirements Engineering Conference, RE 2019, Jeju Island, Korea (South)*, September 23–27, 2019, pp. 109–119, <http://dx.doi.org/10.1109/RE.2019.00022>, URL <https://doi.org/10.1109/RE.2019.00022>.
- [6] A. Lavalle, A. Maté, J. Trujillo, Requirements-driven visualizations for big data analytics: A model-driven approach, in: *Conceptual Modeling - 38th International Conference, ER 2019, Salvador, Brazil, November 4–7, 2019, Proceedings*, Springer, 2019, pp. 78–92, http://dx.doi.org/10.1007/978-3-030-33223-5_8, URL https://doi.org/10.1007/978-3-030-33223-5_8.
- [7] A. Fernandez, A. Bergel, A domain-specific language to visualize software evolution, *Inf. Softw. Technol.* 98 (2018) 118–130, <http://dx.doi.org/10.1016/j.infsof.2018.01.005>, URL <https://doi.org/10.1016/j.infsof.2018.01.005>.
- [8] M. Alshakhouri, J. Buchan, S.G. MacDonell, Synchronised visualisation of software process and product artefacts: Concept, design and prototype implementation, *Inf. Softw. Technol.* 98 (2018) 131–145, <http://dx.doi.org/10.1016/j.infsof.2018.01.008>, URL <https://doi.org/10.1016/j.infsof.2018.01.008>.
- [9] H. Santos, V. Dantas, V. Furtado, P. Pinheiro, D.L. McGuinness, From data to city indicators: A knowledge graph for supporting automatic generation of dashboards, in: *The Semantic Web - 14th International Conference, ESWC, Springer*, 2017, pp. 94–108, http://dx.doi.org/10.1007/978-3-319-58451-5_7, URL https://doi.org/10.1007/978-3-319-58451-5_7.
- [10] M. Kintz, M. Kochanowski, F. Koetter, Creating user-specific business process monitoring dashboards with a model-driven approach, in: *Proceedings of the 5th International Conference on Model-Driven Engineering and Software Development, MODELSWARD*, 2017, pp. 353–361, <http://dx.doi.org/10.5220/0006135203530361>, URL <https://doi.org/10.5220/0006135203530361>.
- [11] A. Vázquez-Ingelmo, F.J. García-Peñalvo, R. Therón, Application of domain engineering to generate customized information dashboards, in: *Learning and Collaboration Technologies. Learning and Teaching - 5th International Conference, LCT, Springer*, 2018, pp. 518–529, http://dx.doi.org/10.1007/978-3-319-91152-6_40, URL https://doi.org/10.1007/978-3-319-91152-6_40.
- [12] K. Börner, *Atlas of Knowledge*, MIT Press Cambridge, MA, 2014.
- [13] S. Madhu Sudhan, J. Chandra, IBA Graph selector algorithm for big data visualization using defence dataset, *Int. J. Sci. Eng. Res.* 4 (3) (2013).
- [14] O. Peña, U. Aguilera, D. López-de-Ipiña, Exploring LOD through metadata extraction and data-driven visualizations, *Program* 50 (3) (2016) 270–287, <http://dx.doi.org/10.1108/PROG-12-2015-0079>, URL <https://doi.org/10.1108/PROG-12-2015-0079>.
- [15] C.C. Gray, W.J. Teahan, D. Perkins, Understanding our analytics: A visualization survey, *J. Learn. Anal.* (2017) in press.
- [16] T.T. Tun, A. Bennaceur, B. Nuseibeh, OASIS: weakening user obligations for security-critical systems, in: *28th IEEE International Requirements Engineering Conference, RE 2020, IEEE*, 2020, pp. 113–124, <http://dx.doi.org/10.1109/RE48521.2020.00023>, URL <https://doi.org/10.1109/RE48521.2020.00023>.
- [17] S. Bresciani, M.J. Eppler, The pitfalls of visual representations: A review and classification of common errors made while designing and interpreting visualizations, *Sage Open* 5 (4) (2015) <http://dx.doi.org/10.1177/2158244015611451>, URL <https://doi.org/10.1177/2158244015611451>.
- [18] A. Pang, C.M. Wittenbrink, S.K. Lodha, Approaches to uncertainty visualization, *Vis. Comput.* 13 (8) (1997) 370–390, <http://dx.doi.org/10.1007/s003710050111>, URL <https://doi.org/10.1007/s003710050111>.
- [19] C.R. Johnson, A.R. Sanderson, A next step: Visualizing errors and uncertainty, *IEEE Comput. Graph. Appl.* 23 (5) (2003) 6–10, <http://dx.doi.org/10.1109/MCG.2003.1231171>, URL <https://doi.org/10.1109/MCG.2003.1231171>.
- [20] M. Aufaure, What's up in business intelligence? A contextual and knowledge-based perspective, in: *Conceptual Modeling - 32th International Conference, ER 2013, Springer*, 2013, pp. 9–18.
- [21] D.A.C. Quartel, W. Engelsman, H. Jonkers, M. van Sinderen, A goal-oriented requirements modelling language for enterprise architecture, in: *Proceedings of the 13th IEEE International Enterprise Distributed Object Computing Conference, EDOC, IEEE*, 2009, pp. 3–13, <http://dx.doi.org/10.1109/EDOC.2009.22>, URL <https://doi.org/10.1109/EDOC.2009.22>.
- [22] M.A. Teruel, A. Maté, E. Navarro, P. González, J.C. Trujillo, The new era of business intelligence applications: Building from a collaborative point of view, *Bus. Inform. Syst. Eng.* 61 (5) (2019) 615–634, <http://dx.doi.org/10.1007/s12599-019-00578-3>, URL <https://doi.org/10.1007/s12599-019-00578-3>.
- [23] F. Dalpiaz, X. Franch, J. Horkoff, Istar 2.0 language guide, *CoRR*, abs/1605.07767, 2016, [arXiv:1605.07767] URL <http://arxiv.org/abs/1605.07767>.
- [24] A. Maté, J. Trujillo, X. Franch, Adding semantic modules to improve goal-oriented analysis of data warehouses using I-star, *J. Syst. Softw.* 88 (2014) 102–111, <http://dx.doi.org/10.1016/j.jss.2013.10.011>, URL <https://doi.org/10.1016/j.jss.2013.10.011>.
- [25] J. Horkoff, E. Yu, Interactive goal model analysis for early requirements engineering, *Requir. Eng.* 21 (1) (2016) 29–61, <http://dx.doi.org/10.1007/s00766-014-0209-8>, URL <https://doi.org/10.1007/s00766-014-0209-8>.
- [26] C. Cares, X. Franch, A metamodelling approach for i* model translations, in: *H. Mouratidis, C. Rolland (Eds.), Advanced Information Systems Engineering - 23rd International Conference, CAISE 2011, London, UK, June 20–24, 2011. Proceedings*, in: *Lecture Notes in Computer Science*, 6741, Springer, 2011, pp. 337–351, http://dx.doi.org/10.1007/978-3-642-21640-4_26, URL https://doi.org/10.1007/978-3-642-21640-4_26.
- [27] M. Asadi, G. Gröner, B. Mohabbati, D. Gasevic, Goal-oriented modeling and verification of feature-oriented product lines, *Softw. Syst. Model.* 15 (1) (2016) 257–279, <http://dx.doi.org/10.1007/s10270-014-0402-8>, URL <https://doi.org/10.1007/s10270-014-0402-8>.
- [28] A. Shi-Nash, D.R. Hardoon, Data analytics and predictive analytics in the era of big data, in: *Internet of Things and Data Analytics Handbook*, Wiley Online Library, 2017, pp. 329–345, <http://dx.doi.org/10.1002/9781119173601.ch19>.
- [29] A. Abela, *Advanced Presentations by Design*, Pfeiffer, 2008.
- [30] R. Marty, *Applied Security Visualization*, Addison-Wesley, 2009.
- [31] K. Börner, *Atlas of Knowledge: Anyone Can Map*, MIT Press, 2015.
- [32] Object Management Group (OMG), MOF 2.0 query/view/transformation specification, 2016, <https://www.omg.org/spec/QVT/1.3/PDF>.
- [33] M. Bostock, Data-driven documents, 2019, URL <https://d3js.org/>.
- [34] Plotly, Dash, 2019, URL <https://plot.ly/>.
- [35] Police Department Incident Reports, 2020, <https://data.sfgov.org/Public-Safety/Police-Department-Incident-Reports-2018-to-Present/wg3w-h783>. (Accessed 15 November 2020).
- [36] San Francisco Open Data portal, 2020, <https://datasf.org/opendata/>. (Accessed 15 November 2020).
- [37] A. Lavalle, A. Maté, J. Trujillo, M.A. Teruel, S. Rizzi, Experimental materials, *Zenodo*, 2021, <http://dx.doi.org/10.5281/zenodo.4637123>, URL <https://doi.org/10.5281/zenodo.4637123>.
- [38] C. Wohlin, P. Runeson, M. Höst, M.C. Ohlsson, B. Regnell, A. Wesslén, *Experimentation in Software Engineering*, Springer Science & Business Media, 2012.
- [39] L. Márquez Carpintero, A. Lavalle, A. Maté, J. Trujillo, Grafik – Herramienta para la generación automática de visualizaciones de dashboards en Data Analytics, 2020, URL <http://hdl.handle.net/10045/109458>.
- [40] B.J. Winer, *Statistical Principles in Experimental Design*, McGraw-Hill Book Company, 1962.
- [41] N.R. Goodman, Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction), *The Ann. Math. Statist.* 34 (1) (1963) 152–177.