

# An Intelligent Visual Big Data Analytics Framework for Supporting Interactive Exploration and Visualization of Big OLAP Cubes

Carlos Ordonez  
University of Houston  
USA  
carlos@uh.edu

Zhibo Chen  
Haliburton  
USA  
zchen6@cs.uh.edu

Alfredo Cuzzocrea\*  
University of Calabria & LORIA  
Italy & France  
alfredo.cuzzocrea@unical.it

Javier Garcia-Garcia  
C3, UNAM  
Mexico  
javgar@unam.mx

**Abstract**—*Visual big data analytics* aims at supporting *big data analytics* via visual metaphors, with a plethora of applications in modern settings and scenarios. In all these domains, visual big data analytics paradigms offer several advantages, among which some noticeable ones are: (i) *fast knowledge understanding from big data sets*; (ii) *pattern and trend discovery from big data sets*; (iii) *entity and model discovery from big data sets*; (iv) *sharing insights among organizations*. Among several proposals, *OLAP-based visual big data analytics methodologies and tools* represents a successful case of visual big data analytics frameworks, which is entirely based on OLAP analysis. In this context, an OLAP cube is typically explored with multiple aggregations selecting different subsets of cube dimensions to analyze trends or to discover unexpected results. Unfortunately, such analytic process is generally manual and fails to statistically explain results. On the basis of these considerations, in this paper we propose an innovative OLAP-shaped visual big data analytics framework that incorporates a state-of-the-art statistical technique for supporting exploration and visualization of OLAP data cubes. An experimental evaluation with a medical data set presents statistically significant results and interactive visualizations, which link risk factors and degree of disease.

**Index Terms**—OLAP Cubes, Visual Big Data Analytics, Big Data Analysis, Intelligent Big Data Systems

## I. INTRODUCTION

Nowadays, a relevant line of research is represented by the so-called *visual big data analytics* (e.g., [1], [11], [13]), whose main goal is that of supporting *big data analytics* (e.g., [28], [29], [38]) via visual metaphors. Visual big data analytics is now characterized by a wide spectrum of applicative settings and scenarios, ranging from smart cities (e.g., [10]) to spatio-temporal big data analysis (e.g., [37]), from healthcare (e.g., [21]) to *e-governance* (e.g., [36]), from network simulation and analysis (e.g., [32]) to end-user analytics (e.g., [18]), and so forth.

In all these domains, visual big data analytics paradigms offer several advantages, among which we recall:

- *fast knowledge understanding from big data sets*: thanks to various visual analytics tools, very-large organizations

are allowed to visualize big data sets in a very intuitive way, as to easily (and quickly) deriving useful and “actionable” knowledge from these data sets, for decision making purposes;

- *pattern and trend discovery from big data sets*: very-large organizations can easily discover patterns and trends hidden in the big data sets they produce themselves, thus achieving a substantial competitive advantage on markets, by highlighting those organization variables that have a greater impact on product quality, customer switching, predictive analytics, and so forth;
- *entity and model discovery from big data sets*: thanks to big data visualization, very-large organizations can discover entities and models into big data sets, and, above all, possible correlations among them – this definitively allows very-large organizations to better identify and isolate those organization’s areas that are more relevant for the achievement of the organization’s goals;
- *sharing insights among organizations*: visual big data analytics also supports insight sharing among organizations – after the knowledge discovery phase, organizations are allowed to share these findings in a visual (and intuitive) manner with other organizations, as to improve the degree of collaborative work and the effectiveness of communication.

Among several proposals, *OLAP-based visual big data analytics methodologies and tools* (e.g., [6], [9], [19], [20]) represents a successful case of visual big data analytics frameworks, which is entirely based on OLAP analysis [3], [4], [16]. OLAP is an important set of techniques for exploratory database analysis. In OLAP a large data set is analyzed with multiple aggregations to find interesting results. Such aggregations, computed based on multiple dimension combinations, resemble a multidimensional cube whose mathematical structure is represented by a lattice. In general, cube computations return simple descriptive statistics such as sums, row counts and averages. In this work, we show parametric statistical tests can help analyze the cube with high statistical reliability. On the other hand, we study how to visualize interesting results

\*This research has been made in the context of the Excellence Chair in Computer Engineering – Big Data Management and Analytics at LORIA, Nancy, France

discovered from the cube. The latter are important results to be embedded into a suitable OLAP-based visual big data analytics framework.

From a data mining perspective, statistical tests [33] offer important advantages over statistical models. There are weak assumptions on the probabilistic distribution behind attributes from the data set. For the statistical test we used, a numeric attribute is assumed to have a Gaussian distribution. Statistical tests can produce statistically reliable results with large or small data sets, whereas most data mining models generally require large data sets to find significant results. Statistical tests are based on simple equations that can be efficiently evaluated with SQL queries because they generally do not require vector or matrix manipulation. By themselves, statistical tests generally require multiple trial and error runs before significant findings are produced. Moreover, each such run requires users to vary multiple parameters or select different dimension subsets from the data sets.

In this work, we propose to combine dimension lattice traversal and parametric statistical tests to identify significant metric differences between cube cells. We present a 2D interactive visualization of the OLAP cube based on a checkerboard that enables isolating and interpreting significant measure differences between two similar cuboids, which differ in one dimension and have the same values on the remaining dimensions. Cube exploration and visualization is performed by automatically generated SQL queries. The latter are important results to be embedded into a suitable OLAP-based visual big data analytics framework.

Based on such issues, we developed algorithms that automate this process of exploring and analyzing a data set by combining OLAP cubes with parametric statistical tests. By enhancing standard exploratory OLAP techniques with statistical tests, we are able to prove that our findings are indeed significant, as opposed to obtaining findings from simple number comparisons. Our algorithms automatically analyze all cuboids from a multidimensional cube while applying statistical tests to discover significant differences in cube measure attributes. While our current application are medical databases, the algorithm enables the analysis of any OLAP database in which the goal is to find specific sets of dimensions that significantly change measure values. This problem is computationally challenging because the cube dimension lattice, which forms the foundation on which the statistical tests are based, represents a combinatorial search space. Due to this, we also introduce several algorithmic and systems optimizations that work towards improving the performance of this exhaustive comparison process.

The paper is organized as follows. Section II describes in details a reference case study that fixes and makes clearer the goals of our research. Basic definitions for OLAP databases and statistical tests, as well as an example are introduced in Section III. The process by which we apply statistical tests on all the cuboids of a dimensional cube using OLAP is explained in Section IV. Section V explains our research on the visual analysis of the cube. Section VI discusses related

work. presents conclusions and directions for future research. This paper significantly extends the previous workshop paper [27], where the basic technique has been firstly introduced.

## II. REFERENCE CASE STUDY

Figure 1 shows a reference case study for the visual OLAP framework proposed in this research.

The case study focuses on a *Smart City Environment*, which is at now of relevant interest for visual big data applications (e.g., [5], [31]). Here, a suitable *Multidimensional Model* is built in order to capture the target analytics goals. From this multidimensional (analytics) model, several OLAP data cubes are derived, as storing the aggregations of interest for the specific application scenario (e.g., averages of cars in a certain city road, counts of citizens in a certain square, and so forth). These data cubes populate the so-called *OLAP Data Cube Layer*. Directly interfaced with the latter layer, our *Visual OLAP Framework* builds the proper exploration and visualization algorithms over target data cubes.

The framework then is integrated with a conventional *Big Data Architecture* as to support the common big data processing tasks. Finally, a proper *Visual Big Data Analytics Framework* is in charge of supporting the visual big data analytics activities (including suitable end-user interfaces), as to provide several *OLAP-Shaped Visual Big Data Analytics Tools* supporting a variety of visual decision-making processes (e.g., [22]).

In this so-delineated context, *privacy of OLAP cubes* (e.g., [8]) as well as performance over Cloud scenarios (e.g., [7]) are other serious concerns to be considered, but left as future work.

## III. DEFINITIONS

We now provide basic OLAP cube definitions. The input data is  $F$ , a fact table containing  $n$  records having  $d$  cube dimensions [16],  $D = \{D_1, \dots, D_d\}$  and a set of  $e$  measure [16] attributes  $A = \{A_1, A_2, \dots, A_e\}$ . The mathematical structure representing all subsets of dimensions and their containment with each other is called the dimension lattice [16]. In OLAP cube processing, the fundamental idea is to compute aggregations (sum(), count()) on measures  $A_i$  by subsets of dimensions (i.e. cuboids or cuboids)  $G$  s.t.  $G \subseteq D$ , effectively performing aggregations at varying granularity levels. The set of all potential aggregations at a certain level in the lattice (i.e. with a fixed number of dimensions) is called a cuboid and one specific aggregation group is called a cell.

In our algorithms, aggregations are used to derive univariate statistics such as  $\mu, \sigma$ , which in turn are the basic input elements in the equations of a parametric statistical test. Section IV explains in more detail cube exploration, the parametric statistical test and how they are combined.

*Example:* Figure 2 presents a simple example of a cube having three dimensions  $D_1, D_2, D_3$  and one measure  $A_1$ . Each face of the cube represents a 2-dimensional cuboid. In this example, there are two sets of cell pairs within one cuboid that differ in exactly one dimension. The difference in fill

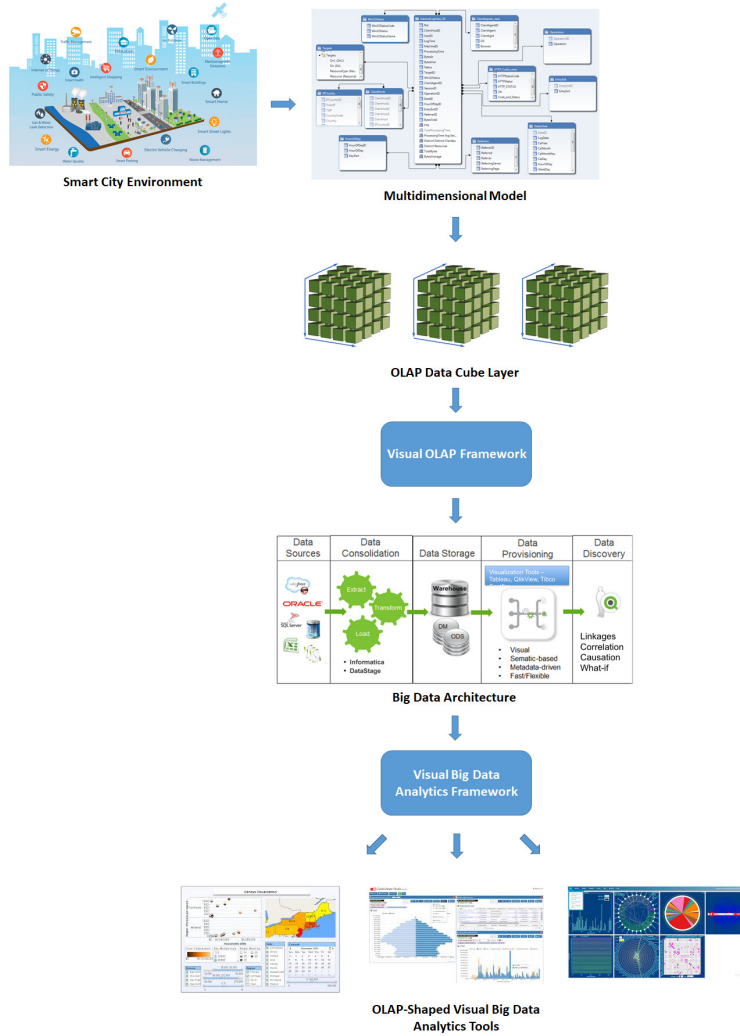


Fig. 1. A reference case study: visual OLAP over a smart city environment.

pattern indicates there is a significant difference on a specific measure attribute  $A_1$ .




Fig. 2. Discovering pairs of subsets with significant measure differences.

## IV. INTEGRATING OLAP CUBES AND STATISTICAL TESTS

### A. Statistical Tests

Prior work [12], [14], [15] generally focused on looking for interesting or unusual patterns on single cells of the cuboids. As such, these works perform simple comparisons between the individual cells or between cells in different aggregations to obtain results. In contrast, we use a parametric statistical test to compare the population means [33] of pairs of groups. Our approach exhibits the following advantages: Two large groups of any size can be compared including groups with very different number of elements (e.g. a large and a small group). The means comparison test takes into account data variance, which measures overlap between the corresponding pair of populations. In the case of OLAP, dimensions can be used to focus on highly similar groups, differing in a few dimensions. It represents a natural extension of OLAP computations since it relies on distributive aggregations [14]. Cube measures are assumed to have a normal distribution, which is a reasonable assumption in most cases.

Let us now describe the means comparison parametric test in more formal terms. This statistical test compares the means,  $\mu_1$  and  $\mu_2$ , from two similar but independent populations of sizes  $N_1$  and  $N_2$ , respectively. We use a null hypothesis  $H_0$  of  $\mu_1 = \mu_2$  with the goal of finding pairs of populations in which  $H_0$  can be rejected. In order to reject  $H_0$ , we aim to reach a high reliability (confidence) value  $1 - p$ , where  $p$  often calls in the following thresholds ( $p \leq \{0.01, 0.05, 0.10\}$ ). When  $H_0$  can be rejected with high confidence, we are able to accept the alternative hypothesis,  $H_1$ , which asserts the complementary comparison  $\mu_1 \neq \mu_2$ . Our parametric statistical test uses a two-tailed test which allows finding a significant difference on both tails of the Gaussian distribution. In order to determine the lowest possible  $p$ , we first determine the random variable  $z$  with a standard probabilistic distribution  $N(0, 1)$  based on Equation 1:

$$z = \frac{|\mu_1 - \mu_2|}{\sqrt{\sigma_1^2/N_1 + \sigma_2^2/N_2}}, \quad (1)$$

where  $\mu_i, \sigma_i$  correspond to the estimated mean and standard deviation for groups (populations) 1, 2, respectively. When  $N$  is large for both groups, it is only necessary to compare the  $z$  value with  $z_{p/2}$  in the cumulative probability table for  $N(0, 1)$ . When either group or both groups are small, we need to compute an additional value, the degrees of freedom:  $df = N_1 + N_2 - 2$ , which together with  $z$ , is used to lookup the  $p$  value on the t-student distribution table.

*Performing multiple statistical tests on cuboids:* We have two main objectives for the application of parametric statistical test: (1) Discovering significant differences between two groups in a cuboid on at least one measure. We should mention that finding two or more significant measure differences based on the same dimension combination is desirable, but unusual. (2) When there exists a significant difference we isolate those groups that differ in one dimension, which can explain a cause-effect relationship. Even though dimensions are assumed independent the aggregation automatically groups records with correlated dimensions together. Therefore, if a high correlation exists in binary dimensions it will be automatically considered.

With respect to goal (1), when applying a statistical test a significant difference can only be supported by a small  $p$ -value which takes into account both the means and the standard deviation of the distributions. The smaller the  $p$ -value the more likely the difference in the cube measure value between both groups is significant. It is expected that many measure differences will not be significant, making the search problem on the dimension lattice expensive. With respect to goal (2), the algorithm aims to discover significant differences in highly similar cube cells because that helps point out which specific dimension “triggers” a significant change on the cuboid measure. In other words, finding a significant measure difference, between two highly dissimilar groups, makes a cause-effect explanation difficult, since such difference may be caused by the interaction of two or more cube dimensions. However, those less significant measure differences can be stored on additional tiers.

## B. Exploration and Visualization Algorithm

We introduce an algorithm that integrates cube exploration, statistical tests and visualization. This algorithm extends our previous algorithm [25] with visualization and interactive exploration. Our algorithm has the following goals: (1) exploring all cuboids from  $F$  (when  $d \leq 10$ ). Otherwise, exploring all cuboids based on  $k$  dimensions manually selected by the user s.t.  $k < d$ ; (2) performing the statistical test for every cube pair; (3) selecting significant pairs differing in  $\delta = 1$  cube dimensions; (4) interactive visual exploration of the cube, together with statistically significant results; (5) efficient visualization of associated image data per cube cell.

Our algorithm basically computes the entire cube, exploring the entire dimension lattice and then applies statistical tests for every pair. The algorithm assumes a low  $d$  or alternatively low  $k$ , binary dimensions, which is common in medical databases. Our tool applies a top-down approach exploring all cuboids from a cube, working level-wise. Further details can be found in [25].

## V. VISUAL ANALYSIS OF THE CUBE

In this section, we explain interactive exploration and visualization of a medical set with our algorithms. First, we explain our computer and DBMS setup and the input medical data set used as a fact table. Second, we explain lattice exploration guided by visualization of cuboids at different dimensionalities. Third, we explain how we visualize significant cuboid cell pairs as well as associated attributes, including image data.

### A. Computer Setup and Data Set

Our system was developed in Java that automatically generates SQL queries. JDBC is used to connect to the DBMS. We used the SQL Server DBMS, running on a computer with 3.2GHz CPU, 4GB of RAM, and 1TB disk.

We performed experiments on a real data set coming from the medical domain. Our medical data set contains profiles of  $n = 655$  patients and has 25 attributes containing categorical, numeric and image data. There were medical measurements such as weight, heart rate, blood pressure and pre-existence of related diseases. Finally, the data set contains the degree of artery narrowing (stenosis) for the four heart arteries. All numeric attributes were converted to binary dimensions. There were  $d = 12$  binary dimensions (e.g. gender, hypertension Y/N),  $e = 4$  measures (artery disease measurement)  $f = 9$  image attributes representing a standardized image of the heart.

We now explain parameter settings. For our medical data sets our goal was to explore the entire dimension lattice. Therefore, we used all  $d$  dimensions. The settings for parameters were as follows.  $p = 0.01$ ,  $\delta = 1$ , which can be interpreted as follows. We want to find significant measure differences, with 99% confidence, on all group pairs differing in one dimension. A group pair in the cube can have from 1 to  $d$  dimensions, out of which one will be different. It is possible, but unlikely, that a group pair has significant differences in two or more measures.

## B. Application in Medicine

Table I shows actual significant findings on the medical data set, introduced above. Each row represents the comparison between two patient groups, differing in one dimension indicated by “0/1”. Cube dimensions are risk factors for heart disease including family history of heart disease, diabetes, gender, high cholesterol and high blood pressure. For a given group, each matching dimension will be “0” indicating absence of a risk factor, “1” indicating presence of a risk factor, or “All” when such dimension was ignored in the aggregation. When  $p$  is small it indicates two highly similar groups of patients, differing in exactly one risk factor have a highly different degree of disease (artery narrowing).

## C. Visualizing Cuboids

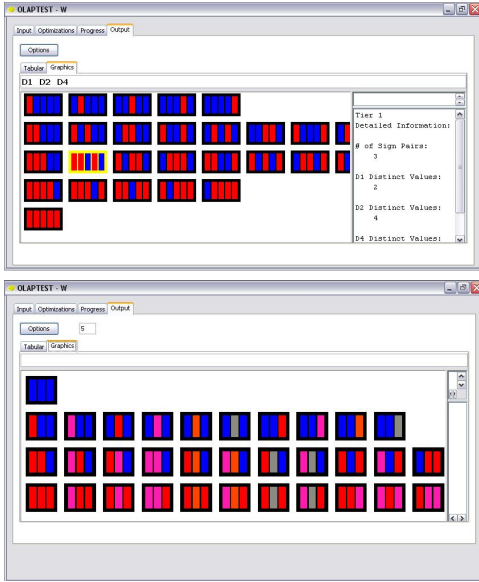


Fig. 3. Visualization of cuboids.

Our visualization architecture can be separated into two main objectives: visualization of cuboids and visualization of significant cell pairs.

In order to visualize the cuboids and allow for quick navigation between different cuboids, we created a two-tiered design, as shown in Figure 3. In this design, the first-tier represents a lattice-like display that shows the various combinations of the  $k$  dimensions through the use of blocks of items. Each item represents one specific dimension and can be colored red, to represent on, or blue, to represent off. A set of  $k$  items represents one possible combination of the chosen dimensions and is called a block. Since OLAP is an exhaustive process that traverses all possible combinations of dimensions, the total number of blocks is  $2^k$ . We chose to arrange the blocks by the number of dimensions that are in the cuboid because it allows for easy navigation. In addition to just viewing the blocks, additional information, such as the total number of significant pairs, is also provided with mouse navigation. Since the  $k$  dimensions are inputted by the user prior to the execution of

the algorithm, this first-tier does not require any retrieval of data from the database.

Once a block is selected, the cuboid that it represents will be displayed to show the second tier of our design. In this level, as shown in Figure 3, on the lower screenshot, there still exists block and items. However, in this case, each block represents one specific combination of the selected dimensions in the cuboid, while each item represents a specific distinct value of the dimension. We use a different color to represent each distinct value for a certain dimension. As such, a red item in one dimension would represent a different value than a red item for another dimension. We provide a collapsible legend that provides information on the values that are represented by each color for each dimension. Note that this more detailed level not only contains the specific values for the selected cuboid, but also shows those blocks that form a significant pair with a block in the chosen cuboid. Once again, placing the mouse over a block will bring up additional information on the panel to the right. Selecting a specific block in this tier will bring up statistical data regarding this population, such as the mean and standard deviation. Additionally, a sample of actual records in the original data set that belong to this block can be viewed. In order to generate the visualization of the specific cuboid, we perform a single retrieval on the final result set with the application of appropriate predicates to reduce the size of the returned values. Since the final result table is indexed on all dimensions, we can quickly retrieve the required records. From this set, we can determine which distinct dimensions were involved in a significant cell pair.

## D. Visualization of Significant Cell Pairs



Fig. 4. Visualization of Significant Cell Pairs.

The second portion of our visualization goal is to easily view and navigate the significant pairs returned by the algorithm. In our application, significant results are represented by a green line linking two blocks of values together, as can be seen in Figure 4. The thickness of the line represents the amount of significance between the two blocks. Should more than one measure be significant between any two blocks, then the thickness would also be altered to reflect this situation. As with the previous visualizations, placing the mouse over the line would provide additional information, such as  $p$ -value. Selecting a line would provide the user with detailed information regarding the overlap between the Gaussians of

TABLE I  
MEDICAL DATA SET: GROUP PAIRS WITH SIGNIFICANT MEASURE DIFFERENCES.

$D_1$ FamHist	$D_2$ Diab	$D_3$ Gender	$D_4$ HighChol	$D_5$ highBP	$N_1$	$N_2$	$A_1$ LAD	$A_2$ RCA
0	All	All	1	0/1	35	23	$p > 0.1$	$p \in [0.01 - 0.05]$
0/1	All	All	1	0	35	26	$p > 0.1$	$p < 0.01$
All	0/1	All	All	All	47	157	$p < 0.01$	$p < 0.01$

the two blocks, or populations, as seen in Figure 4. In this case, the panel to the right would display a graph showing two Gaussians, each representing one of the blocks. The user can thus graphically view the overlap present with the notion being that the smaller the amount of overlap, the greater the significance of the result. Extra information, such as mean and standard deviation, would also be provided on the right panel. In addition, we also provide a feature that allows for the selection of multiple significant lines. In this case, the right panel would display multiple sets of Gaussians to allow the user to visually discern the difference in overlap between different pairs. The creation of these connecting lines requires a single-pass on the final result set. We apply a filter on this result set to only return those significant pairs whose dimensions fit within the selected cuboid. Since the DBMS is optimized for fast retrieval, the time to obtain and generate this visualization is also very efficient. Once the data is obtained, they are stored in either main memory or, if the results are too large, are stored in a binary file for quick access. In this way, the Gaussian curves that are shown when a line is selected can be quickly displayed since the required information is already in RAM memory.

## VI. RELATED WORK

Cube exploration is a well researched topic. OLAP and a classification of aggregations originates in the seminal paper [14]. In [15] the authors put forward the plan of creating smaller, indexed summary tables from the original large input table to speed up aggregating executions. In [30] the authors explore tools to guide the user to interesting regions in order to highlight anomalous behavior while exploring large OLAP data cubes. This is done by identifying exceptions, that is, values in cells of a data cube that are significantly different from the value anticipated, based on a statistical model. In contrast, we propose to use statistical tests to do pair-wise comparison of neighboring cells in cuboids to discover significant metric differences between similar groups. We identify such differences giving statistical evidence about the validity of findings.

There has been research on visualizing OLAP cubes. Recent work can also be found on visually and interactively exploring data warehouses. The authors for [35] explore the requirements for analyzing a spatial database with an OLAP tool. This work shows the need to apply spatial data techniques, used in geographic information systems, for OLAP exploration, in which drill up/down, pivoting, and slicing and dicing provide a complementary perspective. In contrast, our work relies on

statistical tests to explore OLAP cubes and can automatically detect significant metric differences between highly similar groups. Additional visualization work was completed in [23] where the mapping of the Cube Presentation Model, a display model for OLAP screens, involves visualization techniques from the Human-Computer Interaction field. In [17], the author presents a rigorous multidimensional visualization methodology for visualizing n-dimensional geometry and its applications to visual and automatic knowledge discovery. The application of visual knowledge discovery techniques is possible by transforming the problem of searching for multivariate relations among the variables into a two-dimensional pattern recognition problem. A framework for exploration of OLAP data with user-defined dynamic hierarchical visualizations is presented in [34]. Even though this study emphasizes the use of visualization tools to explore data warehouses, we propose a tool that not only gives the user visual aids to explore the data, but also to present the user with a novel method of highlighting interesting features of the cubes by means of statistical tests. Indexing is one method of increasing performance in the searching of images and the authors in [2] proposed a multilevel index structure that can efficiently handle queries on video data. Our work is related to applying data mining in medical data sets to improve heart disease diagnosis [24], [26].

## VII. CONCLUSIONS

Within the wider context of a OLAP-based visual big data analytics framework, we presented an innovative system that combines the exploratory power of OLAP cubes with the statistical reliability of statistical tests. Cube exploration is used to automatically analyze all subsets of dimensions, freeing the user from having to manually set parameters. Pairs of similar groups are isolated and compared with a statistical test to discover specific pairs that cause a significant difference in some measure value. Such differences represent statistically reliable results. The OLAP cube is depicted using a two-tier design that allows the user to quickly switch between cuboids. Significant results are visually shown using connecting lines and also provide additional information with one or more Gaussian graphs that graphically show the overlap between two record populations. We presented an application in the medical domain to improve heart disease diagnosis.

Cube visualization is a fertile research topic since it involves understanding a multidimensional data set. A 2D cube representation is easier to manipulate than a 3D display, but we would like to compare its strengths and weaknesses with a

3D visualization. We need to study mathematical relationships between the dimensions lattice and cube visual representations. The visualization of cubes requires further study when confronted with a cube having a large number of dimensions. Finally, we want to integrate other statistical tests or statistical models with OLAP cubes.

#### ACKNOWLEDGMENTS

This research has been partially supported by the French PIA project “Lorraine Universit dExcellence”, reference ANR-15- IDEX-04-LUE.

#### REFERENCES

- [1] Nikos Bikakis, George Papastefanatos, and Olga Papaemmanouil. Big data exploration, visualization and analytics. *Big Data Res.*, 18, 2019.
- [2] L. Chen, M. Ozsu, and V. Oria. Mindex: An efficient index structure for salient-object-based queries in video databases. *Multimedia Syst.*, 10(1):56–71, 2004.
- [3] Z. Chen and C. Ordonez. Efficient OLAP with UDFs. In *Proc. ACM DOLAP Workshop*, pages 41–48, 2008.
- [4] Z. Chen, C. Ordonez, and C. Garcia-Alvarado. Fast and dynamic OLAP exploration using UDFs. In *Proc. ACM SIGMOD Conference*, pages 1087–1090, 2009.
- [5] Qimei Cui, Yingze Wang, Kwang-Cheng Chen, Wei Ni, I-Cheng Lin, Xiaofeng Tao, and Ping Zhang. Big data analytics and network calculus enabling intelligent management of autonomous vehicles in a smart city. *IEEE Internet Things J.*, 6(2):2021–2034, 2019.
- [6] Alfredo Cuzzocrea and Svetlana Mansmann. OLAP visualization. In *Encyclopedia of Data Warehousing and Mining, Second Edition (4 Volumes)*, pages 1439–1446. IGI Global, 2009.
- [7] Alfredo Cuzzocrea, Rim Moussa, and Guandong Xu. Olap\*: Effectively and efficiently supporting parallel OLAP over big data. In *Model and Data Engineering - Third International Conference, MEDI 2013, Amantea, Italy, September 25-27, 2013. Proceedings*, pages 38–49. Springer, 2013.
- [8] Alfredo Cuzzocrea and Vincenzo Russo. Privacy preserving OLAP and OLAP security. In *Encyclopedia of Data Warehousing and Mining, Second Edition (4 Volumes)*, pages 1575–1581. IGI Global, 2009.
- [9] Alfredo Cuzzocrea, Domenico Sacca, and Paolo Serafino. Semantics-aware advanced OLAP visualization of multidimensional data cubes. *IJDWM*, 3(4):1–30, 2007.
- [10] Xiaoliang Fan, Baoqin He, and Patrick Brézillon. Context-aware big data analytics and visualization for city-wide traffic accidents. In *Modeling and Using Context - 10th International and Interdisciplinary Conference, CONTEXT 2017, Paris, France, June 20-23, 2017, Proceedings*, pages 395–405. Springer, 2017.
- [11] Mingchen Feng, Jiangbin Zheng, Jinchang Ren, Amir Hussain, Xiuxiu Li, Yue Xi, and Qiaoyuan Liu. Big data analytics and mining for effective visualization and trends forecasting of crime data. *IEEE Access*, 7:106111–106123, 2019.
- [12] Lixin Fu and J. Hammer. Cubist: a new algorithm for improving the performance of ad-hoc OLAP queries. In *Proc. ACM DOLAP Workshop*, 2000.
- [13] Matteo Golfarelli and Stefano Rizzi. A model-driven approach to automate data visualization in big data analytics. *Information Visualization*, 19(1), 2020.
- [14] J. Gray, A. Bosworth, A. Layman, and H. Pirahesh. Data cube: A relational aggregation operator generalizing group-by, cross-tab and sub-total. In *ICDE Conference*, pages 152–159, 1996.
- [15] H. Gupta, V. Harinarayan, A. Rajaraman, and J.D. Ullman. Index selection for OLAP. In *IEEE ICDE Conference*, 1997.
- [16] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann, San Francisco, 1st edition, 2001.
- [17] A. Inselberg. Visualization and knowledge discovery for high dimensional data. In *UIDIS*, pages 5–24, 2001.
- [18] Hourieh Khalajzadeh, Mohamed Abdelrazek, John Grundy, John G. Hosking, and Qiang He. Bidaml: A suite of visual languages for supporting end-user data analytics. In *2019 IEEE International Congress on Big Data, BigData Congress 2019, Milan, Italy, July 8-13, 2019*, pages 93–97. IEEE, 2019.
- [19] Sébastien Lafon, Fatma Bouali, Christiane Guinot, and Gilles Venturini. Hierarchical reorganization of dimensions in OLAP visualizations. *IEEE Trans. Vis. Comput. Graph.*, 19(11):1833–1845, 2013.
- [20] Luca Leonardi, Gerasimos Marketos, Elias Frenzos, Nikos Giatrakos, Salvatore Orlando, Nikos Pelekis, Alessandra Raffaetà, Alessandro Roncato, Claudio Silvestri, and Yannis Theodoridis. T-warehouse: Visual OLAP analysis on trajectory data. In *Proceedings of the 26th International Conference on Data Engineering, ICDE 2010, March 1-6, 2010, Long Beach, California, USA*, pages 1141–1144. IEEE Computer Society, 2010.
- [21] Carson K. Leung, Yibin Zhang, Calvin S. H. Hoi, Joglas Souza, and Bryan H. Wodi. Big data analysis and services: Visualization on smart data to support healthcare analytics. In *2019 iThings/GreenCom/CPSCoM/SmartData International Conference, Atlanta, GA, USA, July 14-17, 2019*, pages 1261–1268. IEEE, 2019.
- [22] Audrone Lupeikiene, Viktor Medvedev, Olga Kurasova, Albertas Caplinskis, and Gintautas Dzemyda. Decision support via big multidimensional data visualization. In *Joint Proceedings of the BIR 2015 Workshops and Doctoral Consortium, Tartu, Estonia, August 26-28, 2015*, pages 1–8. CEUR-WS.org, 2015.
- [23] A. S. Maniatis, P. Vassiliadis, S. Skiadopoulos, and Y. Vassiliou. Advanced visualization for OLAP. In *Proc. ACM DOLAP Workshop*, pages 9–16, New York, NY, USA, 2003. ACM Press.
- [24] C. Ordonez. Association rule discovery with the train and test approach for heart disease prediction. *IEEE Transactions on Information Technology in Biomedicine (TITB)*, 10(2):334–343, 2006.
- [25] C. Ordonez and Z. Chen. Evaluating statistical tests on OLAP cubes to compare degree of disease. *IEEE Transactions on Information Technology in Biomedicine (TITB)*, 13(5):756–765, 2009.
- [26] C. Ordonez, N. Ezquerro, and C.A. Santana. Constraining and summarizing association rules in medical data. *Knowledge and Information Systems (KAIS)*, 9(3):259–283, 2006.
- [27] Carlos Ordonez, Zhibo Chen, and Javier García-García. Interactive exploration and visualization of OLAP cubes. In *DOLAP 2011, ACM 14th International Workshop on Data Warehousing and OLAP, Glasgow, United Kingdom, October 28, 2011, Proceedings*, pages 83–88. ACM, 2011.
- [28] Wullianallur Raghupathi and Viju Raghupathi. Big data analytics in healthcare: promise and potential. *Health Inf. Sci. Syst.*, 2(1):3, 2014.
- [29] P. Russom. Big data analytics. *TDWI Best Practices Report*, Fourth Quarter, 2011.
- [30] S. Sarawagi, R. Agrawal, and N. Megiddo. Discovery-driven exploration of OLAP data cubes. In *EDBT*, pages 168–182. Springer-Verlag, 1998.
- [31] Kamran Soomro, Muhammad Nasir Mumtaz Bhutta, Zaheer Khan, and Muhammad Atif Tahir. Smart city big data analytics: An advanced review. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.*, 9(5), 2019.
- [32] Simon Su, Vincent Perry, and Venkat R. Dasari. Hybrid 2d and 3d visual analytics of network simulation data. In *2019 IEEE International Conference on Big Data (Big Data), Los Angeles, CA, USA, December 9-12, 2019*, pages 3992–3999. IEEE, 2019.
- [33] M. Triola. *Essentials of Statistics*. Addison Wesley, 2nd edition, 2005.
- [34] S. Vinnik and F. Mansmann. From analysis to interactive exploration: Building visual hierarchies from OLAP cubes. In *EDBT*, pages 496–514, 2006.
- [35] A. Voß, V. Hernandez, H. Voß, and S. Scheider. Interactive visual exploration of multidimensional data: Requirements for CommonGIS with OLAP. In *DEXA Workshops*, pages 883–887, 2004.
- [36] Hsiao-Fang Yang, Chia-Hou Kay Chen, and Kuei-Ling Belinda Chen. Using big data analytics and visualization to create iot-enabled science park smart governance platform. In *HCI in Business, Government and Organizations. Information Systems and Analytics - 6th International Conference, HCIBGO 2019, Orlando, FL, USA, July 26-31, 2019, Proceedings, Part II*, pages 459–472. Springer, 2019.
- [37] Lina Yu, Michael L. Rilee, Yu Pan, Feiyu Zhu, Kwo-Sen Kuo, and Hongfeng Yu. Visual analytics with unparalleled variety scaling for big earth data. In *2017 IEEE International Conference on Big Data, BigData 2017, Boston, MA, USA, December 11-14, 2017*, pages 514–521. IEEE Computer Society, 2017.
- [38] Paul Zikopoulos and Chris Eaton. *Understanding Big Data: Analytics for Enterprise Class Hadoop and Streaming Data*. McGraw-Hill Osborne Media, 1st edition, 2011.