

Assessing the effectiveness of OLAP exploration approaches

Technical report

Mahfoud Djedaini¹, Pedro Furtado², Nicolas Labroche¹, Patrick Marcel¹, and Verónica Peralta¹

¹ University of Tours, France
`firstname.lastname@univ-tours.fr`
² University of Coimbra, Portugal
`pnf@dei.uc.pt`

Abstract. Supporting interactive database exploration (IDE) is a problem that attracts lots of attention these days. Exploratory OLAP (OLAP: On-Line Analytical Processing) is an important use case where tools support navigation and analysis of the most interesting data, using the best possible perspectives. While many approaches were proposed (like query recommendation, reuse, steering, personalization or unexpected data recommendation), a recurrent problem is how to assess the effectiveness of an exploratory OLAP approach. In this paper we propose a benchmark framework to do so, that relies on an extensible set of user-centric metrics that relate to the main dimensions of exploratory analysis. Namely, we describe how to model and simulate user activity, how to formalize our metrics and how to build exploratory tasks to properly evaluate IDE systems under test (SUT). To the best of our knowledge, this is the first proposal of such a benchmark. Experiments are two-fold: first we evaluate the benchmark protocol and metrics based on synthetic SUTs whose behavior is well known. Second, we concentrate on two different recent SUTs from IDE literature that are evaluated and compared with our benchmark. Finally, potential extensions to produce an industry-strength benchmark are discussed in the conclusion.

1 Introduction

While there exist a set of benchmarks recognized by the database community as relevant for evaluation and comparison of performance of database systems, such as the benchmarks from TPC organization, there is yet no commonly agreed upon benchmark for evaluating to what extent database systems help the user during data exploration. Roughly speaking, TPC benchmarks assess data retrieval, and not data exploration. However, supporting exploration of databases is of prime importance, especially in a context of big, distributed and heterogeneous data, as shown in a recent survey of the topic [17]. Both researchers and companies that supply data analysis tools are increasingly focused on mechanisms for improving user experience, in particular aids for effective data exploration. This is partly because better data exploration experiences will give a

competitive edge to the companies selling those tools. As researchers and companies implement, test and tune alike their Interactive Data Exploration (IDE) solutions, a major issue they face is how to assess and compare solutions, improvements and alternatives.

In this work we focus on the context of OLAP analysis of data, as an important use case of IDE. OLAP is defined as the process of analyzing multidimensional datasets (cubes), online, interactively, summarizing key performance indicators (called measures) from different perspectives or axes of analysis (called dimensions). For example, consider sales of products over time at different shops. Dimensions include product, time and shop, and measures include the amount of sales. Dimension hierarchies are another important concept as they allow analyzing data at different levels of detail. For instance, we might analyze sales by product, brand or type of product, taking advantage of a product hierarchy, study weekly, montly and yearly evolution, according to a time hierarchy, and compare shops, locations, and zones, depending on the way shops are hierarchically organized. Most frequent OLAP queries consist of aggregating measures at specific hierarchical levels, and then comparing measure values, analyzing the impact of specific dimensions, and explaining, somehow, the visualized values. Both beginner and proficient users need to navigate the cube in order to understand the data and to accomplish their analysis tasks. We call this process, an interactive data exploration (IDE) approach. All kinds of data navigation aids, such as query reuse, query steering or query personalization, can be used to suggest next course of actions to users during their exploration.

In order to motivate our work, let's now consider the following toy example: a user navigating OLAP sales cube faces an unexpected difference between sales in year 2014 and year 2015 for a product P in France. The user will then explore the surrounding region of the cube by mean of OLAP operators such as roll-up (at the Europe level for example), drill-down (at the month level for example) and slices (for other products) to find evidences that may explain and corroborate the first fact. The use might even get some support from a system that automatically proposes next moves in the analysis [2,13]. We consider that the surrounding region of the first interesting fact corresponds to a neighborhood that has to be covered to ensure the exploration task success. If one wants to evaluate this particular data exploration, it is then possible to measure several metrics such as the number of queries that the user needed to cover this neighborhood, the ratio of this area that has been finally discovered, the ratio of the rest of the cube that the user had to visit to reach this result etc. So far, the assessment of data exploration through quality measures has been overlooked by the database community, but we can benefit from experience in the fields of information retrieval and exploratory search [34], which are particularly driven by the quality of the user's experience and metrics for measuring it.

In this paper we propose a benchmark to compare approaches for aiding in exploratory OLAP. The main objective of the benchmark is to provide a framework where a System Under Test (SUT) executes and produces suggestions for next data analysis moves, and those suggestions are evaluated according to

a set of metrics that measure the quality of exploration. In this way, a developer can easily test and evaluate alternatives. The benchmark can be used with any SUT, to evaluate any strategy that one may design. It reports objective measures for a set of metrics that characterize the degree to which the SUT fulfills certain objectives.

The benchmark was designed with a set of guiding principles in mind. It has to be easy to use by anyone, considering that a developer or researcher working in an OLAP exploration tool or algorithm, should be able to quickly plug his approach to the benchmark and use it, without requiring complex development or setting up of schema, dataset or OLAP exploration characteristics. The benchmark should also return objective evaluation metric results that are independent of the approach being tested. This means that both the mechanisms of the benchmark and the evaluation metrics must be agnostic of the IDE approach. We fulfill this objective by making sure that the user and evaluation models of the benchmark are open and published as part of the benchmark specification, not based on some specific proposed exploratory OLAP approach, and by defining metrics that classify the quality of exploratory OLAP using state-of-the-art quality-of-exploration metrics.

So far, the assessment of data exploration through quality measures has been overlooked by the database community, but we can benefit from experience in the fields of information retrieval and exploratory search [34], which are particularly driven by the quality of the user’s experience and metrics for measuring it. We were able to adapt metrics that include **Engagement and Enjoyment**, **Information Novelty**, **Task Success**, **Task Time** and **Learning and Cognition** to our context. For those identified metrics we defined how they should be implemented and evaluated in exploratory OLAP.

We conduct two kind of experiments in order to validate the benchmark and to compare actual exploratory OLAP approaches. For validation, we create three ”synthetic” SUTs with expected outcomes: a baseline SUT which returns purely random next move suggestions (called ”random”), and a naive SUT, which changes just one attribute in the query (called ”naive”). The comparison of results between random and naive helps validate that the benchmark is able to rank according to the quality of exploration.

Concerning evaluation of existing exploratory OLAP approaches, we analyze two state-of-the-art interactive data exploration approaches, Cinecubes [13], and falso [1], that are two different approaches that were never compared before. The ability to compare and to reach conclusions about these approaches using the benchmark shows the importance of the benchmark. With the benchmark, we can finally characterize and compare the approaches in an independent way.

The paper is organized as follows. Section 2 discusses related work. Section 3 explains how interactive explorations can be scored and defines the benchmark metrics. Section 4 presents the benchmark itself. First it describes how the SUT is interfaced with the benchmark, then it discusses how the benchmark is initialized with data and sessions, and how it simulates the user accomplishing tasks and finding interesting answers to those tasks. Experimental results are

	Category	Input			Output
		DB instance	Query log	Current query	
[6]	Automatic exploration	✓		✓	Tuples
[2]	Automatic exploration		✓	✓	Sequence of queries
[12]	Automatic exploration	✓	✓	✓	Queries
[4]	Automatic exploration		✓	✓	Queries
[13]	Visual optimization	✓		✓	Queries
	Automatic exploration				result highlighting
[14]	Visual optimization			✓	Query
[30]	Data prefetching	✓	✓	✓	Tuples
[29,31]	Data prefetching	✓		✓	Tuples
[18]	Data prefetching	✓		✓	Sequence of queries
[28]	Data prefetching		✓	✓	Queries

Table 1. Interactive cube exploration techniques signatures

discussed in section 5, including both validation of the benchmark and analysis of existing exploratory OLAP approaches. Finally, section 6 concludes the paper with considerations on potential extensions to produce industry-strength benchmark.

2 Related Work

The variety of database exploration approaches Many approaches have recently been developed to support interactive database exploration (IDE), as illustrated by a recent survey of the topic [17]. Techniques range from Visual optimization (like query result reduction [5]), automatic exploration (like query recommendation [9]), assisted query formulation (like data space segmentation [32]), data prefetching (like result diversification [19]) and query approximation [16]. The core of most of these approaches consists of a function that, given the database instance and users’ history with the database (i.e., past and current queries), computes new relevant queries, tuples or visualizations that are meant to support user exploration.

Given the exploratory nature of OLAP analysis of multidimensional data (see e.g., [30,18]), many exploration techniques have been specifically developed in the context of interactive OLAP exploration of data cubes. Table 1 lists these exploration approaches, indicating their categories (in terms of those proposed in [17]), and report their inputs and outputs. For instance, the PROMISE prefetching approach [28], that predicts a query based on a Markov Model constructed out of the server’s log, corresponds to a function with signature $\langle L, \langle q \rangle \rangle \rightarrow \langle q' \rangle$, where L is the query log, q is the current user query and q' is the predicted query.

Measuring the quality of an exploration Measuring the quality of exploration has attracted a lot of attention in Information Retrieval, in particular in the field

of Exploratory Search³ [34] that can be defined as a search paradigm centered on the user and the evolution of their knowledge. It is particularly driven by the quality of the user's experience, and metrics for measuring it have been categorized as follows. **Engagement and Enjoyment** measures the "degree to which users are engaged and are experiencing positive emotions". It includes "the amount of interaction required during exploration", the "extent to which the user is focused on the task". **Task Success** assesses "whether the user reaches a particular target" and finds a "sufficient amount of information and details" along the way. **Information Novelty** measures the "amount of new information encountered". **Task Time** measures the "time spent to reach a state of task completeness". **Learning and Cognition** measures the "attainment of learning outcomes", "the amount of the topic space covered" and "the number of insights acquired". While these categories have been proposed in the context of web search, they make perfect sense for interactive database exploration, and we next focus on measures that have been proposed in the literature in these categories.

User engagement measures are popular in web search to measure how a user is engaged in using a website or search engine. Many implicit measures have been proposed [23] to track online behavior. These measures are classically categorized in activity (how a website is used), loyalty (return of users to a website) and popularity (how much a website is used). While loyalty and popularity essentially make sense for relative comparison of websites, activity enables measuring engagement for a particular website independently of other websites. The most commonly used activity metrics include number of queries per session, number of clicks, number of clicks per query, dwell (presence) time (see e.g., [10,33]).

Task success is well studied in information retrieval, with even conferences devoted specifically to this, like the TREC conference⁴. Task success is traditionally measured with precision/recall-based measures, which supposes that the target of the task is known. In this case, roughly speaking, recall measures how complete the answer to a query is, while precision measures how noisy the answer to a query is.

Many works have been interested with measuring information novelty in relational databases. For instance, in [11], the authors propose to describe the data space covered by a session with a vector of the tuples accessed by the queries of the session. In [24], the authors propose the notion of access area to capture the portion of the dataset a user is interested in. In [19], the authors use a similar notion to propose query result diversification. In data cubes exploration, Sarawagi [30] assimilates novelty with the most informative constraints so that the expected distribution of a cube's cell values - based on a maximum entropy principle - is closer to the actual observed values. Here, a constraint is defined at an aggregate level of the observed cells and is expressed as a sum over the values of a subset of the observed cells. It is then expected that bringing more constraints modifies the expected distribution of values and thus allows to reduce

³ <http://wp.sigmod.org/?p=1183>

⁴ <http://trec.nist.gov/>

the divergence between the observation and the expectation. The constraints that best reduce this divergence is declared to be the most informative.

Measuring task time may seem straightforward, but one needs to carefully define what is timed and how to report it. Performance related metrics like query per hour can be adapted from TPC benchmarks to this end.

Finally, measuring learning and cognition has attracted lots of attention in learner models [8]. Learner models are central components of intelligent tutoring systems, that infer a student’s latent skills and knowledge from observed data. A very influential and widespread accepted model is the Knowledge Tracing model [7]. Knowledge tracing is a Bayesian network allowing to measure the probability that a skill is mastered when resolving a problem (opportunity to use the skill). The model relies on four parameters, usually experimentally tuned: $P(L_0)$: the probability the skill is already mastered before the first problem, $P(T)$: the probability the skill will be learned at each opportunity to use the skill (transition from not mastered to mastered), g : the probability the resolution is correct if the skill is not mastered (guess), s : the probability a mistake is made if the skill is mastered (slip). The probability that the skill L at opportunity n is mastered is the probability the skill is learned at step $n - 1$ or not learned at step $n - 1$ but learned at this step n . It can be computed as: $P(L_n|X_n = x_n) = P(L_{n-1}|X_n = x_n) + (1 - P(L_{n-1}|X_n = x_n)) \times P(T)$ where:

$$P(L_{n-1}|X_n = 1) = \frac{P(L_{n-1})(1-s)}{P(L_{n-1})(1-s) + (1-P(L_{n-1}))g}$$

$$P(L_{n-1}|X_n = 0) = \frac{P(L_{n-1})s}{P(L_{n-1})s + (1-P(L_{n-1}))(1-g)}$$

$X_n = 1$ (resp. 0) means problem n has been correctly (resp. not correctly) resolved.

Current benchmarks for decision support, big data and analytic workloads TPC proposes a number of popular benchmarks and metrics for assessing the performance of database systems, covering time, performance, price, availability or energy consumption (see Table 2). However, while TPC acknowledges the importance of the explorative nature of decision support queries (see e.g., the OLAP interactive queries in the TPC-DS benchmark), none of the existing TPC metrics are appropriate for measuring database exploration support in the sense of the categories proposed in Exploratory Search. A recent benchmark targets analytical workloads [22], but it too overlooks assessing the quality of interactive data exploration by proposing metrics covering only query response time, tuning overhead, data arrival to query time, storage size and monetary cost.

OLAP-specific literature for the benchmark implementation Interestingly, the literature on OLAP already provides the building blocks for benchmarking cube exploration. OLAP has been the subject of specific benchmarks, like the TPC-H-based Star Schema Benchmark (SSB) [25]. SSB models a realistic use case of sales analysis, for which realistic instances with skewed data can be produced with the PDGF data generator [26]. Realistic OLAP workloads can be generated with the CubeLoad session generator [27]. CubeLoad takes as input a cube schema and creates the desired number of sessions according to templates mod-

Metrics	TPC-H	TPC-DS	TPC-DI	TPCx-HS	TPCx-BB
Query per hour	✓	✓		✓	
Price/performance	✓	✓	✓	✓	✓
Availability date	✓	✓	✓	✓	✓
Power/performance	✓	✓		✓	✓
Power	✓				
Throughput	✓	✓	✓	✓	
Load time		✓			
Power test elapsed time		✓			

Table 2. Metrics of relevant TPC benchmarks

eling various user exploration patterns: users with limited OLAP skills pursuing a specific analysis goal, more advanced users navigating with a sequence of slice and drill operations, users tracking unexpected results with explorative sessions. OLAP literature also provides techniques for characterizing analytic behaviors [28,4]. In these works, the user’s behavior is defined as a Markov model, whose states are built from the past queries of the user, and the transitions between states are weighted by the probability of observing a query after another in the user’s query log. Finally, OLAP literature also provides characterizations of interesting data in the multidimensional space. Discovery-driven analysis of data cube [29,30,31,6] aimed at measuring potentially surprising data, knowing already evaluated queries. These work characterize surprising data as being groups of tuples that are connected (usually one OLAP operation apart), and that, taken altogether, appear to be meaningful (usually unexpected, in the sense of e.g., information theoretic measures).

3 Evaluating an exploration

This section precisely describes how interactive explorations can be scored, by implementing the metrics related to user experience identified in the previous section. We first start with presenting formally the concepts needed to define an exploration in an OLAP context.

3.1 Exploration in an OLAP context

Our benchmark incorporates the explorative and interactive nature of IDE by considering user sessions as first class citizens. We define a session as a sequence of queries over a datawarehouse and a log as a set of sessions. In what follows, a log can be associated to one particular user profile (representing this user’s activity) or can represent the overall activity (being the union of all user logs).

Definition 1 (Session and log). *A session $s = \langle q_1, \dots, q_k \rangle$ of length $|s| = k$ is a sequence of k OLAP queries over a multidimensional schema. A log $L = \{s_1, \dots, s_p\}$ is a finite set of sessions.*

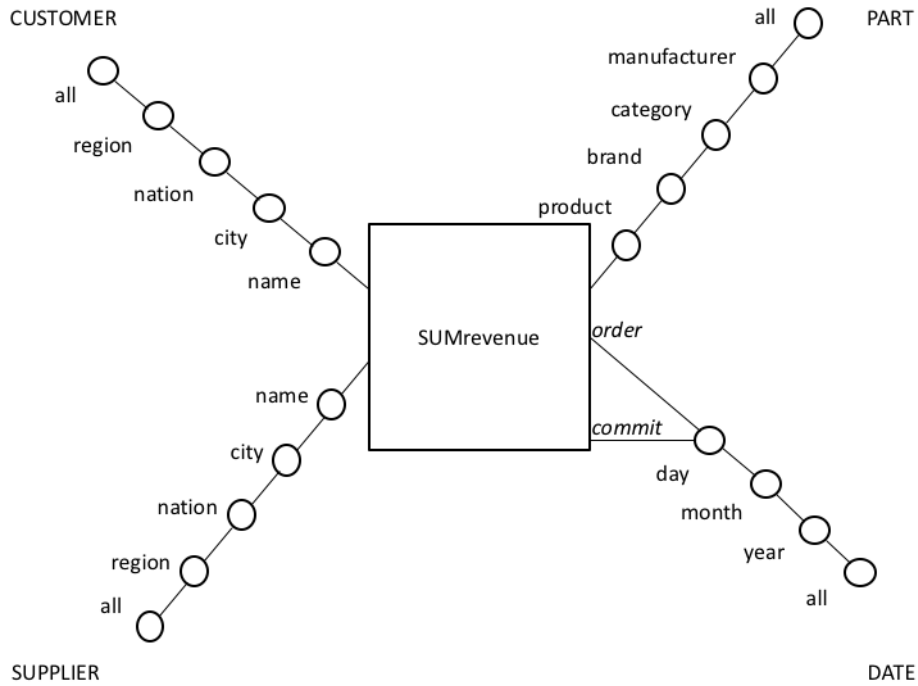


Fig. 1. Multidimensional schema used in the examples

The OLAP queries we consider are dimensional aggregate queries over a data cube [15]. To keep the formalism simple, we consider cubes under a ROLAP perspective, described by a star schema [20]. We consider that a dimension consists of one hierarchy, and we consider simple hierarchies without branches, i.e., consisting of chains of levels. Formally, a cube schema consists of i) n hierarchies, each hierarchy being a set $Lev(h_i) = \{L_0, \dots, L_d\}$ of levels together with a *roll-up* total order \succeq_{h_i} of $Lev(h_i)$, ii) a *measure* attribute m , defined on a numerical domain $Dom(m)$. For a n -dimensional cube, a *group-by set* is an element of $Lev(h_1) \times \dots \times Lev(h_n)$.

More precisely, let \mathcal{L} be a set of attributes called levels, and for $L \in \mathcal{L}$, a member is an element of $Dom(L)$. A hierarchy h_i is a set $Lev(h_i) = \{L_0, \dots, L_d\}$ of levels together with a *roll-up* total order \succeq_{h_i} of $Lev(h_i)$, which is such that, for any L_j and L_k in $Lev(h_i)$, $L_k \succeq_{h_i} L_j$ if $Rollup(L_j) = L_k$.

A *multidimensional schema* (or, briefly, a *schema*) is a triple $\mathcal{M} = \langle A, H, M \rangle$ where:

- A is a finite set of *levels*, whose domains are assumed pairwise disjoint,
- $H = \{h_1, \dots, h_n\}$ is a finite set of *hierarchies*, (such that the $Lev(h_i)$'s for $i \in \{1, \dots, n\}$ define a partition of A);
- a finite set of *measure* attributes M , each defined on a numerical domain $Dom(m)$. In what follows we consider that M is a singleton.

Example 1. Figure 1 displays the conceptual model for a multidimensional schema that will be used in the following examples. This schema is based on that of the SSB [25]. It is formalized by the triple $\mathcal{SSB} = \langle A_S, H_S, M_S \rangle$ where:

- $A_S = \{name^c, city^c, nation^c, \dots\}$ is the set of all *levels*,
- $H = \{h_{CUSTOMER}, \dots, h_{orderDATE}\}$ is the set of *hierarchies*, such that, for instance, $nation^c \succeq_{h_{CUSTOMER}} city^c$
- $M_S = \{SUMrevenue\}$, is the measure.

Definition 2 (OLAP query). A query q over a n -dimensional cube schema is a pair $\langle g, P \rangle$ where g is a group by set and $P = \{p_1, \dots, p_n\}$ is a set of Boolean predicates, one for each hierarchy, whose conjunction defines the selection predicate for q ; they are of the form $L = v$, or $L \in V$, with L a level, v a value, V a set of values. Conventionally, we note $p_i = TRUE_i$ if no selection on h_i is made in q (all values being selected);

Example 2. Consider schema $\mathcal{SSB} = \langle A_S, H_S, M_S \rangle$. A query over this schema is $q_1 = \langle \langle nation^c, all^p, year^{od}, all^{cd}, all^s \rangle, \{TRUE^c, TRUE^p, year^{od} = 2014 \text{ or } 2015, TRUE^{cd}, TRUE^s\} \rangle$. It asks for the yearly revenue by customer nations, for year 2014 or 2015.

During their sessions, after each query is processed, users inspect the cube cells retrieved by the query. A cell is an element of a cube that has a particular coordinate (called reference) and measure value. The answer to a query $q = \langle g, P \rangle$, denoted $answer(q)$, is the set of cells whose references are defined by g and P .

Definition 3 (Cells). Let $g = \langle L_1, \dots, L_n \rangle$ be a group by set. A reference $r = \langle m_1, \dots, m_n \rangle$ is an element of $Dom(g) = Dom(L_1) \times \dots \times Dom(L_n)$. A cell is a pair $\langle r, v \rangle$ where r is a reference and v a numerical value in $Dom(m)$.

Example 3. Here are two cells retrieved by query q_1 :

- $c_1 = \langle \langle FRANCE, all, 2015, all, all \rangle, 10000 \rangle$,
- $c_2 = \langle \langle Italy, all, 2015, all, all \rangle, 20000 \rangle$.

Thanks to the the popular OLAP operations (roll-up, drill-down, slice-and-dice), users navigate the cube by exploring cells neighborhood, by querying at coarser granularity (roll-up), finer granularity (drill-down) or reaching siblings in a hierarchy. This is formalized through relations between cells.

Let L_j and L_k be two levels of a hierarchy h such that $L_k \succeq_h L_j$, and let $m_j \in Dom(L_j)$ and $m_k \in Dom(L_k)$. We note $m_k \succeq_h m_j$ if m_j is the ancestor of m_k in hierarchy h . \succeq_h is a transitive relation whose transitive closure is denoted \succeq_h^* . We note $m_k \approx_h m_j$ if m_j is the sibling of m_k (they have the same direct ancestor) in hierarchy h . Given two cells $c = \langle r, v \rangle$ and $c' = \langle r', v' \rangle$, we consider the classical roll-up relation over cells defined by: $c \succeq^* c'$ if, for all hierarchies h_i it is either $r(i) = r'(i)$ or $r(i) \succeq_{h_i}^* r'(i)$. The sibling relation between cells, noted \approx , is defined as $c \approx c'$ if there exist a hierarchy h_i for which $r(i) \approx_{h_i} r'(i)$ while for all other hierarchies $j \neq i$ it is $r(j) = r'(j)$.

Example 4. Consider cells $c_1 = \langle \langle \text{FRANCE}, \text{all}, 2015, \text{all}, \text{all} \rangle, 10000 \rangle$,
 $c_2 = \langle \langle \text{Italy}, \text{all}, 2015, \text{all}, \text{all} \rangle, 20000 \rangle$ from the previous example, and cell
 $c_3 = \langle \langle \text{Europe}, \text{all}, 2015, \text{all}, \text{all} \rangle, 100000 \rangle$. It is $c_1 \approx c_2$ and $c_1 \succeq c_3$.

Definition 4 (Neighborhood of a cell). *The rollup neighborhood of a cell c is the set of all cells c' such that $c' \succeq_c c$. The drilldown neighborhood of a cell c is the set of all cells c' such that $c \succeq_c c'$. The sibling neighborhood of a cell c is the set of all cells c' such that $c \approx c'$. The OLAP neighborhood of a cell c is the union of its rollup neighborhood, its drilldown neighborhood and its sibling neighborhood.*

The neighborhood of a group of cells C , noted $neighborhood(C)$ is the union of the neighborhoods of each cell of the group. Intuitively, the neighborhood of a group of cells defines a zone of the cube to be explored to analyze this group of cell.

A user is represented by a log, i.e., the user’s past explorations. This allows to characterize a user’s behavior by constructing a generative model, in the spirit of what has been successfully applied in OLAP for data prefetching [28]. We now give the formal definition the generative model representing a user’s behavior.

Definition 5 (Generative model). *Let L be a set of sessions characterizing a user. The generative model to represent this user’s behavior is a Markov model of order one, i.e., a graph $\langle S, P \rangle$ where S is the set of queries of L and $P : S \times S \rightarrow [0, 1]$ denotes the probability function for the state transition, computed as $P(q_1, q_2) = \frac{sessions(\langle q_1, q_2 \rangle)}{sessions(q_1)}$ where q_1 and q_2 are queries and $sessions(s)$ gives the number of sessions where the sequence s appears.*

Definition 6 (User). *Let S be a set of sessions and x be a percentage. A user u_x is a tuple $u_x = \langle s^{log}, s^{seed}, g \rangle$ where $S = s^{log} \cup s^{seed}$, $|s^{log}| = x \times |S|$ and g is the generative model built from s^{log} .*

Finally, a task for a session consists of a set of cells to be analyzed by a user. This set of cells is given under the form of a session, i.e., consists of the cells retrieved by the queries of this session.

Definition 7 (Task). *Let s be a session and u be a user. A task is a pair $\langle s, u \rangle$.*

3.2 Metrics

We now present the metrics used to score an OLAP exploration. As explained in Section 2, we adopt a categorization of the metrics proposed in the field of Exploratory Search [34]. For each category, we propose a primary metric and a secondary metric, with the idea that secondary metrics can be used to counter-balance primary ones. Note that metrics of different categories have been defined so that the overlapping between them is minimal: User engagement relates only to the number of queries, novelty to cells, task success to cell neighborhood and task time only to time. Only Learning and cognition overlaps with novelty task

success since it aims at measuring the skill of finding new and relevant information. In what follows, let $u = \langle s^{log}, s^{seed}, g \rangle$ be a user, let $t = \langle s^0, u \rangle$ be a task for user u and let $s = \langle q_1, \dots, q_k \rangle$ be a session produced for the resolution of a task t .

User engagement and enjoyment We use two popular and simple activity metrics used in web search: click depth as primary metric, to represent overall activity, and number of clicks per query to represent how focused this activity is. Dwell time, another popular activity metric, better fits in the Task time category. In the web search context, a click correspond to following a hyperlink (i.e., an HHTP query). In the context of the benchmark, a click corresponds to a new query. The metrics are defined as follows:

- **Query depth (QD, primary)** = k , i.e., the number of queries in the session, needed for resolving a task.
- **Focus (F, secondary)** = $\frac{\max(\{|focus(q)| | q \in s\})}{|s|}$, where $focus(q) = \langle q, \dots, q' \rangle \subseteq S$ such that for all $q_i, q_{i+1} \in focus(q)$ the cells retrieved by q_{i+1} are in one of the neighborhood of the cells retrieved by q_i . Intuitively, this is to measure for a query q , the length of the chain of queries starting from q that are successively distant of only one OLAP operation.

Information Novelty Capturing user interest in the data explored can be done by measuring the access area [24]. In our context, this access area would be the set of tuples (recorded in a fact table) contributing to form the cells of a query result. As this area corresponds to tuples that are not actually presented as answers to queries (since, being an OLAP context, these tuples are aggregated), data of interest is better captured with view area, i.e., the cells presented in the answers. This is defined by: given a set of query $Q = \{q_1, \dots, q_n\}$, the view area of Q is $va(Q) = \bigcup_{q \in Q} answer(q)$.

In a view area, not all data is interesting in the sense that it brings novel knowledge.

We measure interestingness degree as a simple normalized entropy:

$$interest(C) = \left(- \sum_{i=1}^m p(i) \log(p(i)) \right) / \log(m) \quad (1)$$

with $|C| = m$, $C(i)$ is the i^{th} value of the set C and $p(i) = \frac{C(i)}{\sum_{i=1}^m C(i)}$ denotes the estimated value of cell i or its probability.

The primary metric then quantifies the amount of interesting data found in the session. The secondary metric measures the increase in view area compared to the user’s log view area.

The metrics are defined as follows:

- **Relevant new information (RNI, primary)** = $\max_{q \in s} (interest(va(q)))$.
- **Increase in view area (IVA, secondary)** = $\frac{|va(s) \setminus va(s^{log})|}{|va(s) \cup va(s^{log})|}$,

Task Success Intuitively, a task consists of investigating what can be said of a group of cells C coming from a task $\langle s, u \rangle$. The extent to which a task is complete consists of assessing how much of the neighborhood of this group of cells has been retrieved during the resolution of the task. A simple way of measuring it is with recall and precision. Recall is the primary metrics since consistently with exploratory search, we consider OLAP navigation as a recall oriented activity (what matters most is to minimize the number of false negative). The metrics are defined as follows, for a group of cells C :

- **Recall (R, primary)** = $\frac{|va(s) \cap neighborhood(C)|}{|neighborhood(c)|}$.
- **Precision (P, secondary)** = $\frac{|va(s) \cap neighborhood(C)|}{|va(s)|}$

Task Time Measuring task time is done by adapting metrics of existing TPC benchmarks. We need to measure the time for the SUT to produce its output and to process the queries needed for the resolution of the task. The primary metric comes from the TPC-DS benchmark and measures the number of queries per the time taken to resolve the task. The secondary metric simply measures the task elapsed time. The metrics are defined as follows:

- **Query per seconds (QpS, primary)** = $\frac{k}{\sqrt{T_o \times T_e}}$, where T_o is the overall time for the SUT to produce its outputs and T_e is the overall query execution time.
- **Task elapsed time (TET, secondary)** = $T_o + T_e$, where T_o is the overall time for the SUT to produce its outputs and T_e is the overall query execution time.

Learning and cognition We adapt the Knowledge Tracing model to our context. Recall that knowledge tracing is a Bayesian network for measuring the probability that a skill is mastered when resolving a problem. In our context, the skill assessed is the ability of finding interesting information in the neighborhood of a group of cells C , $neighborhood(C)$.

The problem to solve (the opportunity to test the skill) is the creation of a new query, i.e., an opportunity to find more interesting cells in $neighborhood(C)$. $X_n = 1$ if the n^{th} query finds at least one more unknown cell of this neighborhood where novelty increases for those cells compared to query $n - 1$. It is 0 otherwise.

The challenge is then to define the four parameters of the knowledge tracing. In our context, it makes sense to set these parameters based on the generative model since this model represents the past of the user. Intuitively, the skill improves if the queries of the past sessions retrieve cells that are novel and in the neighborhood of C . Parameters $P(L_0)$, g and s are set as follows.

- $P(L_0)$ is the average novelty of the cells in $neighborhood(C)$ retrieved by the generative model, weighted by the fraction of cells of this neighborhood in the generative model.
- $P(g)$ is the probability that novelty increases for the cells in $neighborhood(C)$, from one query to the next one in the generative model.

- $P(s)$ is the probability that novelty decreases for the cells in $neighborhood(C)$, from one query to the next one in the generative model.
- $P(T)$ is the proportion of queries of the sessions for which the novelty of cells in the neighborhood of C increases compared to that of the previous query.

The primary metric is defined as in the classical knowledge tracing model, by: **Learning (L, primary)** = $P(L_n|X_n = x_n) = P(L_{n-1}|X_n = x_n) + (1 - P(L_{n-1}|X_n = x_n)) \times P(T)$ where: $P(L_{n-1}|X_n = 1) = \frac{P(L_{n-1})(1-s)}{P(L_{n-1})(1-s) + (1-P(L_{n-1}))g}$ and $P(L_{n-1}|X_n = 0) = \frac{P(L_{n-1})s}{P(L_{n-1})s + (1-P(L_{n-1}))(1-g)}$.

The secondary metric measures the average progression of the learning curve. It is defined by the geometric mean of the proportional growth of the probabilities: **Learning growth rate (LGR, secondary)** =

$$\prod_{i=1}^m \left(1 + \frac{P(L_i|X_i = x_i) - P(L_{i-1}|X_{i-1} = x_{i-1})}{P(L_{i-1}|X_{i-1} = x_{i-1})} \right)^{1/n} \quad (2)$$

where n is the length of the session.

4 The Benchmark

We define a System Under Test (SUT) as being an Exploratory OLAP system to be evaluated. The main functionality of an exploratory OLAP tool is to support user navigation over the data, to learn some insights from the data. Navigation is defined as a sequence of views that corresponds to a sequence of queries that are submitted to the system. In this context, the sequences the SUT helps to develop are the target of our evaluation. In this section we define the interface between the SUT and the benchmark and how the benchmark runs an experiment.

4.1 Interfacing with a SUT

In order to assess a SUT, the benchmark, simulates a user and interacts with the SUT. The SUT first builds its inner structures, if any, and obtains input metadata from the benchmark. Conceptually, a SUT requires as input all or part of the following parameters: the database (schema and instance), user traces (i.e., sequences of queries collected into the query log) and the active user's current exploration (a sequence of queries). Let D denote the set of all database instances for a given schema, Q denotes the set of all possible queries over this schema, S denotes the set of all sequences of queries (i.e. $Q \times Q \times \dots \times Q$), and 2^A denotes the power-set of a set A . The functionality of a SUT can be defined generically as doing the transformation: $\langle D, 2^S, S \rangle \rightarrow S$. Once the SUT is ready, the evaluation protocol starts resolving a task, successively calling the SUT to provide suggestions.

In practice, the benchmark is a Java program where SUTs can be plugged to be evaluated. Its code and javadoc are available for SUT programmers on

BitBucket⁵. Basically, a SUT is sought twice, (1) before starting the evaluation so it can initialize, and (2) whenever a next move suggestions is requested. From the benchmark point of view, SUTs are only seen as a black boxes that perform what they are asked to perform, through a contract. Practically, a SUT is a class that implements an interface that exposes two functions *readMetadata* and *nextSuggestion*. Function *readMetadata* is called before starting the actual evaluation process, so the SUT can read and initialize its internal structure. Its parameter is a *Metadata* object whose getters allow to access the cube, the list of users, past user traces, etc. Function *nextSuggestion* is called many times during a task resolution. It provides to the SUT a given user and a current exploration (sequence of queries), and asks the SUT to recommend. It is the responsibility of the benchmark to orchestrate the whole process, and to make sure the functions are called with the right arguments.

4.2 How the Benchmark Works

The benchmark process is composed of three components. The first component initializes the benchmark. It generates the context: the database (i.e., the cube), some sequences of queries (i.e., the log), data skews to simulate interesting observations, and creates user profiles. You do not need to run this component if you reuse an existing context, but you can also create a new context with different schema or generation parameters.

The second component is responsible for the actual evaluation of a SUT. The evaluation is a simulation of a user’s actual navigation, whereby the benchmark suggests some initial sequence of queries asks the SUT for next move suggestions, then it proposes some continuation, it switches to ask the SUT again, and so on. This allows the benchmark to ask the SUT for suggestions multiple times, in multiple phases and focusing multiple view areas.

The third component is in charge of computing scores and reporting results. It considers the sessions produced with the SUT, and computes values for the quality metrics presented in Section 3.2. The final step is to provide a report aggregating these metrics, the main output of the benchmarking process. **if we do so**

4.3 Component 1: Benchmark Initialization

Initialization consists of the synthesis of an OLAP user environment. The benchmark initialization process is sketched in Algorithm 1 and Figure 2, and consists mainly of data creation and user generation.

Data generation An OLAP database (schema and instance) and a set of user sessions over it are generated in lines 1-3. The default database schema is the one of SSB benchmark [25], but the benchmark can be initialized with any other OLAP schema. We use CubeLoad [27] for automatically generating user

⁵ <https://bitbucket.org/mdjedaini/ea-benchmark>

Algorithm 1 Generating benchmark objects

Input n : number of sessions, m : number of users, s : database scale factor $\langle x_1, \dots, x_m \rangle$: a vector of percentages

Output D : database schema, I : database instance, U : set of users

- 1: $D =$ Generate the database schema
 - 2: $S =$ Generate a set of n sessions over D
 - 3: $I =$ Generate an instance over D according to S
 - 4: $\{s_1, \dots, s_m\} =$ Partition S by Similarity \triangleright user u_i relates to sessions s_i
 - 5: **for each** s_i **do**
 - 6: $(s_i^{log}, s_i^{seed}) =$ Split(s_i, x_i) \triangleright log and seed sessions for u_i
 - 7: Create generative model g_i for s_i^{log} \triangleright abstracts navigation patterns of u_i
 - 8: $u_i = (s_i^{log}, s_i^{seed}, g_i)$
 - 9: $U = \{u_1, \dots, u_m\}$
 - 10: **return** D, I, U
-




Fig. 2. Initialization process: (i) generation of data and user sessions; (ii) creation of user profiles (visible sessions, hidden sessions and generative model).

sessions. CubeLoad generates realistic OLAP workloads, taking as input a cube schema and the desired number of sessions. Its templates enable the creation of a large number of sessions representing varied explorations and patterns. Finally, a realistic database instance is generated with PDGF [26]. We use the more frequent selection predicates in the log of sessions to produce data skew in the most queried zones of the cube.

Users creation While CubeLoad enables the generation of a large workload and creates feasible exploration patterns, it does not assign sessions to specific users. We use an off-the-shelf clustering algorithm [21], using a similarity measure tailored for OLAP sessions [3] to generate "users". In this way a user is characterized by a set of sessions focusing on some zones of the cube (line 4). Each set of sessions is split in two parts: log and seed sessions (line 6). The former constitutes the user log that is exposed to the SUT, so that it can build its own knowledge for suggesting next moves. The latter, not shown to the SUT, is used to seed the benchmark tasks. The size of each user's log is ruled by a parameter. This allows the benchmark to evaluate the SUT when working with novice users versus advanced users, creating tasks with different difficulty levels, in the sense that it is more difficult for a SUT to suggest something interesting to a relatively new user. Finally, a generative model is learned from the log (line 7), inspired

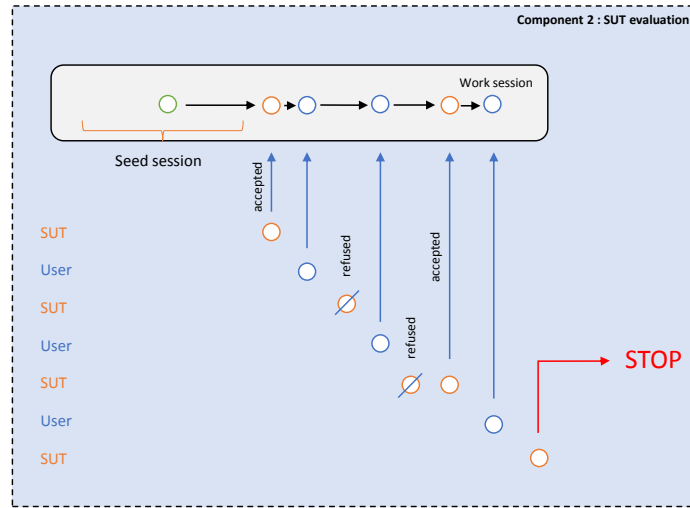


Fig. 3. Task resolution process

by techniques of the OLAP literature [28,4]. This generative model is a Markov Model that is used by Component 2, for simulating the interaction with a user.

4.4 Component 2: Evaluation of a SUT

This component is responsible for the simulation of a navigation, together with the SUT, in order to resolve a given task. A task can be seen as an exercise that has to be solved by SUTs. Tasks are created just before starting a SUT evaluation. The evaluation protocol first provides a seed session, which is a set of seed queries representing part of a navigation, as a context for continuation of the navigation. Then it asks the SUT for a first next move suggestion that consists of one or more queries. After the SUT suggestion, the benchmark decides if it accepts or refuses the suggestion (a real user would either follow the suggestion or not). The probability of discarding the suggestion is given as parameter. The following step is for the benchmark to indicate the next query (a real user may evaluate their own queries). This is chosen as the closest query in the user model to the current query. This new query is then presented to the SUT to suggest again, and the process continues as such until a stop condition. The simulation is ran for a set of tasks (the number of tasks to run is a user-given parameter), and the whole process is preceded by the definition of tasks to accomplish.

4.5 Component 3: Scoring

All the queries recorded during a task resolution are fed to the scoring component so it can compute a score for the SUT for the given task using the metrics defined in Section 3.2. First, each task receives a score for each metric. In practice, a

metric can be seen as a function that takes as input a task resolution (the queries that were played), and provides as output a number that represents the score of the metric for the given task.

5 Experiments

In this section we describe and report results on the experiments designed to validate the proposed benchmark. A first version of the benchmark application was coded in Java, using PDGF [26], CubeLoad [27] and Fuzzy C-medoids [21], as explained in section 4. The tested SUTs were plugged to the benchmark application using the interface class. Experiments use the default schema (SSB) [25] with a scale factor of 1, a small global log of 50 sessions and 375 queries and 9 users with 50% of seed sessions. We generated 100 tasks for each SUTs to resolve. Tests were conducted on a laptop equipped with an i5-3210M CPU @ 2.50GHz and 8GB of RAM.

5.1 Experimental setup

Validation. In order to test benchmark ranking, we compared three synthetic SUTs that have simple behavior, and then expected results. 'Random', the one having worst strategy, returns purely random next move suggestions. 'Naive' generates queries that are one OLAP operation away from the previous query. It naively tries to stay close from the current query, but still chooses the next move randomly within that neighborhood. 'Cheater' uses 'insider information' in order to return good suggestions taken from the log, including seed sessions. Concretely, it replays queries containing cells from the neighborhood of the ones in the seed session, which should fit the user's needs in terms of task success. The goal of this experiment is to confirm that the benchmark ranks these approaches as expected.

Benchmarking existing approaches. We created an experimental setup to compare the following approaches: CineCube [13] and Falseto [2]. CineCube is a multifaceted approach focusing on building a user-friendly sequence of explanations for the analysts. The approach highlights relevant cells in current views and explores automatically expansion into two one-distance children and two one-distance sibling queries, also summarizing the findings. Falseto is an OLAP session composition tool that implements a recommender system based on collaborative filtering [2]. It features three phases: (i) search the log for sessions that bear some similarity with the one currently being issued by the user; (ii) extract the most relevant subsessions; and (iii) adapt the top-ranked subsession to the current user's session.

5.2 Analysis of Experimental Results

Table 3 shows the benchmark results for the tested SUTs. For each SUT, we report its average score and standard deviation for the 100 tasks, for all the benchmark primary and secondary metrics.

	Engagement		Novelty		Success			Time		Learning	
	QD	F	RNI	IVA	R	P	QpS	TET	L	LGR	
Random	102	0,031	1	0,834	0,104	0,001	0,005	545,88	0,855	0,045	
stdev	0	0,012	0	0,190	0,166	0,002	0,003	1032,08	0,117	0,098	
Naive	102	0,117	1	0,699	0,132	0,003	0,004	1005,52	0,852	0,469	
stdev	0	0,083	0	0,210	0,174	0,004	0,004	1425,54	0,123	0,278	
Cheater	39,50	0,232	1	0,180	0,173	0,195	0,074	285,61	0,778	0,473	
stdev	21,76	0,210	0	0,353	0,186	0,380	0,135	1155,29	0,194	0,290	
Falseto	473,94	0,023	1	0,867	0,444	0,003	0,001	12111,78	0,835	0,476	
stdev	39,95	0,003	0	0,229	0,306	0,002	0,000	6234,34	0,103	0,293	
Cinecube	100,4	0,053	1	0,947	0,388	0,215	0,084	1853,59	0,752	0,505	
stdev	21,97	0,079	0	0,064	0,260	0,397	0,178	4091,07	0,284	0,300	

Table 3. Scores of the SUTs

Validation. Globally, the results allow us to rank 'Cheater' highest, followed by 'Naive' and 'Random' with the poorest performance, as expected. Having access to detailed insider information (seed sessions), 'Cheater' achieved a higher task success in less time, and devising less queries. As a consequence, it is more focused, with a lower query depth. However, as it replays existing queries, increase in view area is lower and it provides less learning, but with a more consistent learning curve. 'Random' propose completely random jumps in the multidimensional space, which is less effective (lower task success) and produces unfocused sessions. As it slowly contributes to task resolution, the stop condition (100 queries) stops its execution. That is why it obtains maximum query depth for all tasks (stdev=0) but needs more execution time. Nevertheless, it randomly explores other cube zones, so consequently increases view area and increases learning. As expected, 'Naive' stays half-way between 'Cheater' and 'Random'. By moving always close to the current query, it was able to stay within relevant regions (so succeeding quite well) and being middle focused, but providing some learning and some novelty. As 'Naive', it executes until the stop condition obtaining maximum query depth. As its has a more complex strategy than 'Random', execution times are higher.

Benchmarking existing approaches Results in Table 3 highlight the differences between Falseto and Cinecube and helps deciding which is best in which case. By definition, Falseto generates longer sessions than Cinecube as reflected by the Query Depth score. Falseto also generates queries that are not only related to the neighborhood of the last queries as Cinecube but that are based on collaborative filtering with user past sessions to recommend next analysis moves. This leads Falseto to have more diverse queries than Cinecube. This is an advantage when it comes to explore the data as shown by Task Success Recall of Falseto that is slightly better than Cinecube, but it comes at the cost of a lower precision, because more part of the cube have been explored that are not related to the objective. Falseto is also much slower than Cinecube as reflected by Time primary and secondary metrics.

6 Conclusion

In this paper we proposed the first benchmark for assessing OLAP exploration approaches. Modern OLAP exploration approaches are expected to suggest next moves to users, but an important question is how to evaluate the quality of such suggestions, and how to compare alternatives. Our benchmark uses state of the art techniques to generate data and user traces, and for its metrics definition. The benchmark is easy to use, requiring the SUT tester to write only a well-defined interface, and classifies the SUT according to a set of user-centric metrics. This is an important advance, since existing benchmarks focus almost exclusively on performance, cost or energy. To validate the approach, we have proved that the benchmark correctly ranks a set of strategies for which the behavior is known.

We plan to make all the details of the benchmark public for anyone to use and improve, and our long-term goal is that it serves as a building block of a more general benchmark for exploratory search over databases in general. We are currently working on turning our proposal into an industry-strength benchmark: we are detailing rules, procedures, reporting procedures and documentation; we are investigating the benchmark robustness and its sensitivity to the data and traces; we are applying the benchmark to rank other existing exploratory approaches, as a way to create a regular use base.

References

1. J. Aligon, K. Bouilil, P. Marcel, and V. Peralta. A holistic approach to OLAP sessions composition: The falseto experience. In *DOLAP 2014*, pages 37–46, 2014.
2. J. Aligon, E. Gallinucci, M. Golfarelli, P. Marcel, and S. Rizzi. A collaborative filtering approach for recommending OLAP sessions. *DSS*, 69:20–30, 2015.
3. J. Aligon, M. Golfarelli, P. Marcel, S. Rizzi, and E. Turrinchia. Similarity measures for olap sessions. *KAIS*, 39(2):463–489, 2014.
4. M.-A. Aufaure, N. Kuchmann-Beauger, P. Marcel, S. Rizzi, and Y. Vanrompay. Predicting your next olap query based on recent analytical sessions. In *DaWaK*, pages 134–145, 2013.
5. L. Battle, M. Stonebraker, and R. Chang. Dynamic reduction of query result sets for interactive visualizaton. In *Int'l Conference on Big Data*, pages 1–8, 2013.
6. V. Cariou, J. Cubillé, C. Derquenne, S. Goutier, F. Guisnel, and H. Klajnmic. Embedded indicators to facilitate the exploration of a data cube. *IJBIDM*, 4(3/4):329–349, 2009.
7. A. T. Corbett and J. R. Anderson. Knowledge tracing: Modelling the acquisition of procedural knowledge. *UMUAI*, 4(4):253–278, 1995.
8. M. C. Desmarais and R. S. J. de Baker. A review of recent advances in learner and skill modeling in intelligent learning environments. *UMUAI*, 22(1-2):9–38, 2012.
9. M. Drosou and E. Pitoura. Ymaldb: exploring relational databases via result-driven recommendations. *VLDB J.*, 22(6):849–874, 2013.
10. A. Drutsa, G. Gusev, and P. Serdyukov. Future user engagement prediction and its application to improve the sensitivity of online experiments. In *WWW*, pages 256–266, 2015.
11. M. Eirinaki, S. Abraham, N. Polyzotis, and N. Shaikh. Querie: Collaborative database exploration. *TKDE*, 26(7):1778–1790, 2014.

12. A. Giacometti, P. Marcel, E. Negre, and A. Soulet. Query recommendations for olap discovery-driven analysis. *IJDWM*, 7(2):1–25, 2011.
13. D. Gkesoulis, P. Vassiliadis, and P. Manousis. Cinecubes: Aiding data workers gain insights from OLAP queries. *IS*, 53:60–86, 2015.
14. M. Golfarelli, S. Rizzi, and P. Biondi. myolap: An approach to express and evaluate OLAP preferences. *TKDE*, 23(7):1050–1064, 2011.
15. J. Gray, S. Chaudhuri, A. Bosworth, A. Layman, D. Reichart, M. Venkatrao, F. Pellow, and H. Pirahesh. Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub totals. *Data Min. Knowl. Discov.*, 1(1):29–53, 1997.
16. J. M. Hellerstein, P. J. Haas, and H. J. Wang. Online aggregation. In *SIGMOD*, pages 171–182, 1997.
17. S. Idreos, O. Papaemmanouil, and S. Chaudhuri. Overview of data exploration techniques. In *SIGMOD*, pages 277–281, 2015.
18. N. Kamat, P. Jayachandran, K. Tunga, and A. Nandi. Distributed and interactive cube exploration. In *ICDE*, pages 472–483, 2014.
19. H. A. Khan, M. A. Sharaf, and A. Albarrak. Divide: efficient diversification for interactive data exploration. In *SSDBM*, pages 15:1–15:12, 2014.
20. R. Kimball. *The Data Warehouse Toolkit: Practical Techniques for Building Dimensional Data Warehouses*. John Wiley, 1996.
21. R. Krishnapuram, A. Joshi, O. Nasraoui, and L. Yi. Low-complexity fuzzy relational clustering algorithms for web mining. *IEEE-FS*, 9:595–607, 2001.
22. J. LeFevre, J. Sankaranarayanan, H. Hacigümüş, J. Tatemura, and N. Polyzotis. Towards a workload for evolutionary analytics. In *DanaC '13*, pages 26–30, 2013.
23. J. Lehmann, M. Lalmas, E. Yom-Tov, and G. Dupret. Models of user engagement. In *UMAP*, pages 164–175, 2012.
24. H. V. Nguyen, K. Böhm, F. Becker, B. Goldman, G. Hinkel, and E. Müller. Identifying user interests within the data space - a case study with skyserver. In *EDBT 2015*, pages 641–652, 2015.
25. P. E. O’Neil, E. J. O’Neil, X. Chen, and S. Revilak. The star schema benchmark and augmented fact table indexing. In *TPCTC*, pages 237–252, 2009.
26. T. Rabl, M. Poess, H. Jacobsen, P. E. O’Neil, and E. J. O’Neil. Variations of the star schema benchmark to test the effects of data skew on query performance. In *ICPE’13*, pages 361–372, 2013.
27. S. Rizzi and E. Gallinucci. Cubeload: A parametric generator of realistic OLAP workloads. In *CAiSE 2014*, pages 610–624, 2014.
28. C. Sapia. Promise: Predicting query behavior to enable predictive caching strategies for olap systems. In *DaWaK*, pages 224–233, 2000.
29. S. Sarawagi. Explaining differences in multidimensional aggregates. In *VLDB*, pages 42–53, 1999.
30. S. Sarawagi. User-adaptive exploration of multidimensional data. In *VLDB*, pages 307–316, 2000.
31. G. Sathe and S. Sarawagi. Intelligent rollups in multidimensional OLAP data. In *VLDB*, pages 531–540, 2001.
32. T. Sellam and M. L. Kersten. Meet charles, big data query advisor. In *CIDR*, 2013.
33. Y. Song, X. Shi, and X. Fu. Evaluating and predicting user engagement change with degraded search relevance. In *WWW*, pages 1213–1224, 2013.
34. R. W. White and R. A. Roth. *Exploratory Search: Beyond the Query-Response Paradigm*. Morgan & Claypool Publishers, 2009.