

QUERYVIS: Logic-based Diagrams help Users Understand Complicated SQL Queries Faster

Aristotelis Leventidis

 Northeastern University
leventidis.a@northeastern.edu

Jiahui Zhang

 Northeastern University
zhang.jiahu@northeastern.edu

Cody Dunne

 Northeastern University
c.dunne@northeastern.edu

Wolfgang Gatterbauer

 Northeastern University
w.gatterbauer@northeastern.edu

H.V. Jagadish

 University of Michigan
jag@umich.edu

Mirek Riedewald

 Northeastern University
m.riedewald@northeastern.edu

ABSTRACT

Understanding the meaning of existing SQL queries is critical for code maintenance and reuse. Yet SQL can be hard to read, even for expert users or the original creator of a query. We conjecture that it is possible to capture the logical intent of queries in *automatically-generated visual diagrams* that can help users understand the meaning of queries faster and more accurately than SQL text alone.

We present initial steps in that direction with visual diagrams that are based on the first-order logic foundation of SQL and can capture the meaning of deeply nested queries. Our diagrams build upon a rich history of diagrammatic reasoning systems in logic and were designed using a large body of human-computer interaction best practices: they are *minimal* in that no visual element is superfluous; they are *unambiguous* in that no two queries with different semantics map to the same visualization; and they *extend* previously existing visual representations of relational schemata and conjunctive queries in a natural way. An experimental evaluation involving 42 users on Amazon Mechanical Turk shows that with only a 2–3 minute static tutorial, participants could interpret queries meaningfully faster with our diagrams than when reading SQL alone. Moreover, we have evidence that our visual diagrams result in participants making fewer errors than with SQL. We believe that more regular exposure to diagrammatic representations of SQL can give rise to a *pattern-based* and thus more intuitive use and re-use of SQL.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD'20, June 14–19, 2020, Portland, OR, USA
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6735-6/20/06...\$15.00

<https://doi.org/10.1145/3318464.3389767>

A full version of this paper with all appendices and supplemental material for the experimental study (stimuli, raw data, and analysis code) are available at <https://osf.io/btszh>

ACM Reference Format:

Aristotelis Leventidis, Jiahui Zhang, Cody Dunne, Wolfgang Gatterbauer, H.V. Jagadish, and Mirek Riedewald. 2020. QUERYVIS: Logic-based Diagrams help Users Understand Complicated SQL Queries Faster. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD'20)*, June 14–19, 2020, Portland, OR, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3318464.3389767>

1 INTRODUCTION

SQL is a powerful query language that has remained popular in an age of rapidly evolving technologies and programming languages. Unfortunately SQL queries are often verbose and involve complex logic constructs. This makes them hard to read to a degree where even SQL experts require considerable time *to understand* a non-trivial query.

While the difficulty of *composing* SQL queries has received much attention, it is often just as important to *read and understand* them correctly. For example, SQL queries may require maintenance as database schema or data properties evolve, or when the analysis goals change. Even the development of new queries can be facilitated by understanding and reusing existing ones. A paradigm, successfully employed in projects such as the Sloan Digital Sky Survey [69], is to begin with a query that is similar to the desired one and then modify it as needed. In fact, several systems have been proposed that let users browse and re-use SQL queries in a large repository, including CQMS [45, 46], SQL QuerIE [5, 18], DBease [52], and SQLshare [40]. The key premise of these systems is that *starting from an existing template* should make it easier to specify an SQL query than starting from scratch. However, in order for users to successfully build upon an existing SQL query, they need to understand it first.

Our goal. Compared to query composition, SQL query interpretation is relatively unexplored. Our goal is to provide an approach that simplifies the process of SQL query

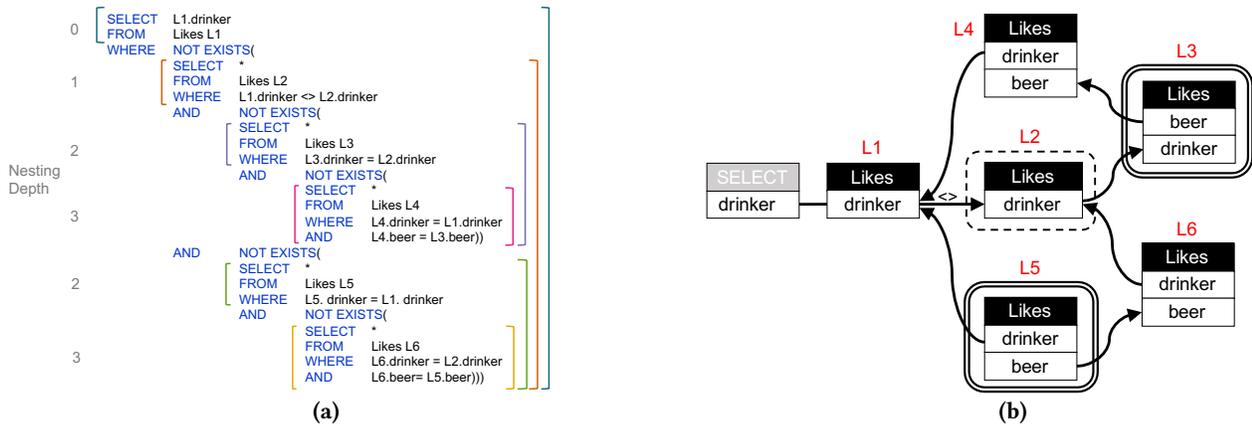


Figure 1: (a): The *unique-set-query* over the bar-drinker-beer schema, whose purpose is to find drinkers that like a unique set of beers. The nesting depth of each subquery is denoted in gray on the left; the scope of each subquery is shown by the brackets on the right and their respective “roots” by the brackets on the left. (b): Our visual diagram for the same query. The red table aliases next to the tables are not part of the diagram and are only placed to illustrate the correspondence to the SQL query. Notice that the visual pattern on the right is the same for different SQL queries that follow the same logical pattern, such as *find beers with a unique set of drinkers* or *find movies with a unique cast of actors* or *find customers with a unique set of purchased items*. Thus our diagrams allow users to inspect and recognize the underlying logical pattern.

interpretation. For this purpose, we propose automatically generated *diagrammatic representations* of SQL queries that capture their logical intent. Our approach is orthogonal to SQL composition and hence can be used to complement any existing SQL development tool, whether visual or not.

Target audience. We target two types of users: Foremost, we like to help users who browse through a repository of existing SQL queries (e.g., their own past queries, or a log of past issued queries over a shared scientific data repository) and try to quickly understand the meaning of such queries. The purpose can be to either find a past query again, or to run queries created by others, or to study existing queries and modify them later. Our paper shows that *diagrams can speed up the process of interpreting existing SQL queries*. The second target is more speculative. We hypothesize (but do not claim to have yet evidence) that providing a formalism for SQL users to help reason in terms of *SQL patterns* can be a helpful process, both while learning SQL and also later when remembering a particular SQL pattern when composing a new query. We provide more detailed illustrating examples for the possibilities in the appendix, available at osf.io/btszh

1.1 Query Interpretation: An Example

SQL queries can be notoriously complex, even when they have a compact description in natural language. By visualizing the logic of an SQL query, we hope to make its logic and intent easier to understand. The following detailed example illustrates this idea.

The unique-set query. Consider the well-known beer drinkers schema by Ullman [73]: Likes(person,beer), Frequents(person,bar), Serves(bar,beer). Suppose we

wish to find *drinkers who like a unique set of beers*, i.e., no other drinker likes the exact same set of beers. The query requires only the Likes table; its SQL text is shown in Fig. 1a. Please take a moment to look at the SQL statement and verify that it correctly expresses the desired query. If this takes you several minutes, it may not be because of your lack of SQL expertise: the logic of SQL is intricate. After some effort, the query can be read as: *return any drinker, s.t. there does not exist any other drinker, s.t. there does not exist any beer liked by that other drinker that is not also liked by the returned drinker and there does not exist any beer liked by the returned drinker that is not also liked by the same other drinker.*

Set theory. From a set logic perspective, this query applies the following logical pattern: Let x be a drinker and $S(x)$ be the set of beers x likes. Our intent is find those x , s.t. there does not exist another drinker $y \neq x$ for which $S(y) \subseteq S(x)$ and $S(y) \supseteq S(x)$. In other words, find *drinkers for which no other drinker has simultaneously a subset and superset of their beer tastes*. Hence, the query “merges” two logical patterns: (1) no other drinker likes a subset of the liked beers, and (2) no other drinker likes a superset of the liked beers. *The other drinker* in both logical patterns must be *the same person*.

First-order logic. SQL and relational calculus are based on first-order logic (FOL). FOL expresses the first pattern (no other drinker likes a subset of the liked beers) as *there does not exist any other drinker $y \neq x$, s.t. all beers liked by y , are also liked by returned drinker x* . Similarly, the second pattern (no other drinker likes a superset of the liked beers) is expressed as *there does not exist any other drinker $y \neq x$, s.t. all beers liked by a returned drinker x , are also liked by y* . Notice that both conditions must be fulfilled *simultaneously* by any

other drinker y , thus our *composite* pattern is a *conjunction* of the two aforementioned patterns, sharing the same other drinker y . Also notice that SQL does not support universal quantification directly. Thus the statement *all beers liked by x are also liked by y* needs to be transformed into the more convoluted *no beer liked by x is not also liked by y* .

Our visual diagrams. Our method is a *diagrammatic representation system* based on FOL that automatically translates such logical patterns from SQL into visual patterns in a way that makes it easier for a user to inspect and recognize them.

Figure 1b shows the visual pattern for the example. A dashed bounding box represents a logical *Not Exists* (\nexists) and a double-lined bounding box represents a *For All* (\forall) quantifier, which are applied to the attributes of the enclosed tables. To read the diagram, we start from its SELECT box and follow the arrows to the next table attribute: The first pattern consists of the set of bounding boxes $L1 \rightarrow L2 \rightarrow L3 \rightarrow L4$ and reads as follows: *Return any drinker (L1), s.t. there does not exist a different drinker (L2), s.t. for all beers liked by the different drinker (L3), they are also liked by the returned drinker (L4)*. The second pattern consists of bounding boxes $L1 \rightarrow L2 \rightarrow L5 \rightarrow L6$, thus sharing the first two boxes.¹ The additional conditions (forming a conjunction with the former) are: *... and s.t. for all beers liked by the returned drinker (L5), they are also liked by the different drinker (L6)*. Reading through the diagram *feels* similar to reading a FOL expression where appropriate symbols identify the predicates and the quantifiers applied to them. This proximity to FOL is not a coincidence, but rather a key feature that preserves and exposes the logic behind SQL queries, yet facilitates their interpretation. Also notice that, in contrast to SQL, we can avoid a double negation and instead use a more intuitive universally quantified statement.

Reading the query visualization may not seem simple at first. Yet notice: (1) any representation system, including classic ER or UML diagrams, may appear cryptic to a novice who sees this representation for the first time; and (2) the logic of the unique-set query is indeed non-trivial. However, because this logic is represented by only a handful of boxes, instead of multiple dense lines of SQL text, it can actually be *easier for readers to recognize*, once the visual conventions become familiar (e.g., after a short tutorial). We will present experimental results verifying this claim later in Section 6.

Common visual patterns. The logical pattern behind a particular query is not unique to the query, and the visual diagram remains the same for queries with identical logical patterns. For example, if we want to find all bars that have a unique set of visitors, the diagram would remain the same

¹The reading order follows a depth-first traversal from the SELECT box with restarts on source nodes: After the path $L1 \rightarrow L2 \rightarrow L3 \rightarrow L4$, the reading starts from L5 (that has no incoming edges) and continues $L5 \rightarrow L6$. Later Section 4.6 has the details.

except for replacing table and attribute names appropriately. This is true even across schemas, e.g., for a query finding all movies with a unique cast in a movie database we obtain the same visual pattern, allowing for the recognition of similarities that are difficult to distill from pure SQL. Thus, our diagrams *expose the underlying logical patterns* to the user in a way that facilitates query interpretation and recall.

1.2 Challenges, Contributions, and Outline

The **challenges** we faced when developing our diagrammatic visualization were to design a representation that (i) can be intuitively learned and quickly understood, (ii) can express a large fragment of SQL, (iii) is not entirely detached from SQL but rather captures the essence of SQL logic, (iv) is minimal in that no visual element is superfluous, (v) is designed based on human-computer interaction best practices, (vi) is unambiguous in that no two queries with different semantics map to the same visualization, and (vii) extends previously existing visual representations of relational schemata and conjunctive queries in a seamless way.

Our **experimental study** with 42 participants shows that existing SQL users can determine the meaning of queries meaningfully faster (-20%, $p < 0.001$) and more accurately (-21%, $p = 0.15$) *using our diagrams alone* instead of standard SQL. There is also some evidence that participants make meaningfully fewer errors (-17%, $p=0.16$) when looking at both *our diagrams together with SQL* instead of SQL alone. These participants were recruited on Amazon Mechanical Turk (AMT) and spent only 2-3 min on a short tutorial with 6 examples of SQL annotated with their respective diagrams. Thus *while the participants had significant prior experience with SQL, they were exposed to our visualizations for only a few minutes and were still faster and typically made fewer errors interpreting the queries*. We can only imagine the improvements if users received more regular exposure to those diagrams and thereby could start to internalize the underlying *logical patterns of SQL queries*. We thus believe that our approach shows a direction that is worthwhile for our community to explore in order to make relational databases more usable [13, 44].

Our main **contributions**, in presentation order, are:

- (1) We survey closely related approaches and explain why visual query builders cannot provide the functionality needed for effective query visualization (Section 2);
- (2) We identify abstract visual design requirements for assisting humans in query understanding (Section 3);
- (3) We present our novel diagrammatic representation of SQL, discuss its origin in first-order logic and diagrammatic reasoning systems, and justify our design choices using a theory of minimal and effective SQL visualizations (Section 4).

- (4) We prove that our diagrams are *unambiguous*, i.e., it is not possible for two different logic representations to lead to the same visual diagram (Section 5);
- (5) We present an empirical validation of our approach with a randomized controlled study involving 42 users, which provides evidence that existing SQL users are meaningfully faster at correctly understanding queries using our diagrams than using SQL text, despite having experienced only minimal prior training on our visualizations (Section 6).

Earlier work and additional material. An earlier vision paper [35] and an interactive system demonstration [25] referred to our approach as *QueryViz*, which we have since renamed to *QueryVis*. Those two short papers described the vision and a prototype implementation, yet lacked a detailed justification of the design, a proof of the diagrams being unambiguous, and an empirical user study. A full version of this paper with all appendices is available as [51]. All supplemental materials for the experimental study (stimuli, raw data, and analysis code) are available at <https://osf.io/mycr2> and on the project web page <http://queryvis.com>

2 RELATED WORK

For decades, SQL has been the main standard for specifying queries over relational databases and this is unlikely to change anytime soon. Thus we do not propose new ways for users to write queries, but instead explore how to help them *understand existing SQL queries*.

Visual query languages. Visual methods for specifying queries have been studied extensively (a 1997 survey by Catarci et al. [14] cites over 150 references) and many commercial database products offer some visual interface for users to write SQL. We focus on the problem of describing and *interpreting a query that has already been written*, which is very different from the problem of helping a user to compose a query. The central difference is that understanding a query requires a focus on the high-level structure, abstracting away low-level details and subtleties [35]. In contrast, to precisely specify a query, possible options and specific details affecting query semantics must be presented. In programming languages, this distinction is clearly made between *visual programming* for developing a program and *program visualization* for analyzing an existing program [58].

Interactive query builders employ visual diagrams that users can manipulate (most often in order to select tables and attributes), while using a *separate query configurator* (similar to QBE’s condition boxes [81]) to specify selection predicates, attributes, and sometimes nesting between queries. dbForge [26] is the most advanced and commercially supported tool we found for interactive query building. Yet it

does not show any visual indication for non-equi joins between tables and the actual filtering values and aggregation functions can only be added in a separate query configurator. Moreover, it has limited support for nested queries: the inner and outer queries are built separately, and the diagram for the inner query is *presented separately and disjointly* from the diagram for the outer query. Thus *no visual depiction of correlated subqueries is possible*. Other graphical SQL editors like SQL Server Management Studio (SSMS) [70], Active Query Builder [3], QueryScope from SQLdep [62], MS Access [54], and PostgreSQL’s pgAdmin3 [61] lack in even more aspects of visual query representations: most do not allow nested queries, none has a single visual element for the logical quantifiers NOT EXISTS or FOR ALL, and all require specifying details of the query in SQL or across several tabbed views *separate from a visual diagram*. DataPlay [1, 2] allows a user to specify their query by interactively modifying a *query tree with quantifiers* and observing changes in the matching/non-matching data. The declared goal is to overcome SQL’s lack of *syntactic locality* (i.e. otherwise similar SQL queries with different quantification can have very different structure). QUERYVIS is designed with the same issue in mind (see [35] and Fig. 2), yet also leverages familiarity with existing visual metaphors for conjunctive queries. In short, current graphical SQL editors *do not provide a single encompassing visualization of a query*. Thus they could not (even in theory) transform a complicated SQL query (such as the one from Section 1.1) into a single visual representation, which is the focus of our work.

Query visualizations attempt to create a visual representation of existing queries. This explicit reverse functionality for SQL has not drawn as much attention as visual query builders. The two projects that come closest in spirit are GraphSQL [15] and Visual SQL [43]. Both are visual query languages that also support query visualization. GraphSQL uses visual metaphors that are different from typical relational schema notations and visualizations, even simple conjunctive queries can look unfamiliar. Visual SQL is closer in design. With its focus on query specification, it maintains the one-to-one correspondence to SQL, and syntactic variants of the same query lead to different representations (see Fig 4 in [35]). The Query Graph Model (QGM) developed for Starburst [39] helps users understand query plans, not query intent. StreamTrace [8] focuses on visualizing temporal queries over streams with workflow diagrams and a timeline. We focus on displaying the underlying logic behind general SQL queries, independent of data.

Information visualization with its goal to help users *understand and analyze data* [19] has recently drawn a lot of attention in the database community. Similar to our objective, visualization researchers help users understand complex relationships, but in data instead of in query logic.

Syntax highlighting. Query editors and clients for major DBMSs have long used *syntax highlighting and aligning* of query blocks and clauses. This is helpful, yet not sufficient to help users understand a query’s intention. In our experiments, SQL queries are auto-indented, and the keywords are capitalized and highlighted in color (see Fig. 3). Still users generally responded better to our approach than to the visually improved SQL text.

Natural language translations. Translating between SQL and NL is an interesting and heavily researched topic, and various ideas are proposed to *explain queries in natural language* [36, 42, 48, 68, 78]. Work in this area convincingly argues that automatically creating effective free-flowing text from queries is difficult and that the overall task is quite different from previous work on creating natural language interfaces to DBMSs. The key limitation of the current state of the art [36, 78] is that it (i) produces long sentences, (ii) is currently limited to simple SQL queries, and that (iii) *textual descriptions do not readily reveal common logical patterns* behind queries. In particular, we are *not aware of any SQL to NL tool available today* that could translate our example from Figure 1 into an intuitive NL representation.

Several papers [2, 17, 59] suggest illustrating the semantics of operators in a data flow program or the semantics of queries by generating **example input and output data**. The result is basically a list of tuples for each relational operator, which, like the natural-language explanation, does not readily reveal the logical pattern behind a query.

3 TASK ANALYSIS & ABSTRACTION

In order to understand the logical pattern behind a query, a user must first perform several low-level tasks. Appropriately determining what these tasks are [27] and abstracting them into domain-independent terms [12] is crucial for designing a visualization that will enable users to achieve their goals. Here we describe our task analysis and abstraction which we will later use to motivate our visualization design.

3.1 User Task Analysis

Users perform tasks over a visualization in order to accomplish their higher-level goal of interpreting an SQL query correctly. To determine the necessary low-level tasks we performed a scenario-based task analysis [27] based on our considerable collective experience using and teaching SQL.

In order to correctly interpret a *conjunctive query with inequalities* (e.g., the one shown in Fig. 3a) the user tasks are:

- (1) **Tables & attributes:** Identify the tables specified in the SQL query and their relevant attributes (i.e., attributes involved in a predicate or selected).

- (2) **Selection predicates:** Identify the attributes of a table that are assigned a selection predicate and the details of that predicate (e.g., $T.attr1 = 4$).

- (3) **Join predicates:** Identify the pairs of attributes with a join predicate as well as the join operator (e.g., $T.attr1 = S.attr1, T.attr2 \geq S.attr2$).

We also wish to support *nested conjunctive queries with inequalities* (we discuss this SQL fragment later in Section 4.4). In order to correctly interpret such nested queries (e.g., the query in Fig. 3b), a user must additionally:

- (4) **Quantifiers:** Identify the logical quantifier (\exists , \neq , or \forall) applied to a set of tables and their attributes.

- (5) **Nesting order:** Identify subqueries, their associated tables, and their the nesting order.

Addressing these additional tasks for nested subqueries adds substantial design challenges.

3.2 User Task Abstraction

Building on our task analysis we conducted a task abstraction [12] in order to discuss these database domain tasks using abstract visualization terms. There are three abstract properties that the visualization should clearly portray: relations, group membership, and hierarchy.

(I) Relations. This property includes user tasks (2) *Selection predicates* and (3) *Join predicates*. The *Selection predicates* task deals with identifying the attribute and the selection predicate. We are thus interested in identifying a relation applied to a single attribute. Similarly, the *Join predicates* task deals with identifying pairs of attributes that join tables and the associated operator for which there is a join condition.

(II) Grouping. This property includes user tasks (1) *Tables & attributes* and (4) *Quantifiers*. The *Tables & attributes* task deals with identifying tables and their relevant attributes that are part of the table. We are thus interested in grouping relevant attributes of a given table as well as distinguishing between tables. The *Quantifiers* task deals with identifying the set of tables to which a logical quantifier is applied. We are thus interested in collecting and grouping all the tables for which the same quantifier is applied.

(III) Hierarchy. The (5) *Nesting order* task deals with identifying the nesting order of subqueries and their associated tables. We are thus interested in identifying the hierarchical structure of the subqueries present in the SQL query. Note that the portrayal of hierarchy is not necessary if we only deal with conjunctive queries as every table will be at the same nesting depth (e.g., the query in Fig. 3a).

4 VISUALIZATION DESIGN

With our task analysis and abstraction established, we now turn to the design of our visualization.

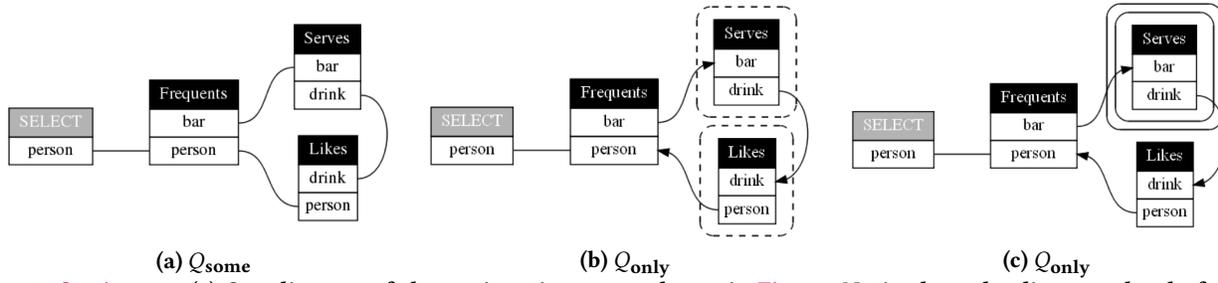


Figure 2: Section 4.8: (a) Our diagram of the conjunctive query shown in Fig. 3a. Notice how the diagram closely follows a familiar relational schema notation. (b) Our diagram of the nested query shown in Fig. 3b. Notice that we have added a dashed bounding box for \exists and the reading order can be found by following the arrows. (c) Fig. 2b can be further simplified through the use of the \forall quantifier (double-lined bounding box), a logical and intuitive operator that does not exist in SQL.

```

SELECT F.person
FROM Frequents F, Likes L, Serves S
WHERE F.person = L.person
AND F.bar = S.bar
AND L.drink = S.drink

```

(a) Q_{some}

```

SELECT F.person
FROM Frequents F
WHERE not exists
(SELECT *
FROM Serves S
WHERE S.bar = F.bar
AND not exists
(SELECT L.drink
FROM Likes L
WHERE L.person = F.person
AND S.drink = L.drink))

```

(b) Q_{only}

Figure 3: (a) Q_{some} : Find persons who frequent some bar that serves SOME drink they like. (b) Q_{only} : Find persons who frequent some bar that serves ONLY drinks they like $\equiv \dots$ some bar that serves NO drink that is NO liked by them.

4.1 Minimal & Effective Visualization

Our goal in this section is to start from two basic primitives of visualization design — *marks* and *channels* [57] — and develop from the ground up, a theory of *minimal* and *effective* SQL visualizations. We intend to demonstrate, using this theory, a more formal way for designing and evaluating visualizations for database systems.

All visualizations are composed of *marks* and *channels* [57]. *Marks* are geometric primitive objects such as points, lines, areas, and volumes. A visual *channel* is a way to control the appearance of a mark, which can be done by changing the position, shape, size, color, texture, orientation, motion path, and many other visual attributes of a given mark. With those basic building blocks for visualizations, we now would like to find visualizations that are minimal and effective:

DEFINITION 4.1. A visualization designed for a set of user tasks is (1) minimal if every mark and channel used is necessary

to accomplish at least one task and (2) effective if every mark and channel used is well-justified for the associated tasks based on human perception studies and visualization best practices.

Note that this definition builds upon previous work in the visualization community. Our definition of effectiveness agrees with Mackinlay [53]. Likewise, our definition of *minimality* combines elements of Mackinlay’s effectiveness and expressiveness criteria [53] and Tufte’s guidance to maximize the data-to-ink ratio [72].

Like Brehmer & Munzner we define *target user tasks* as “domain- and interface-agnostic operations performed by users” [12, 56]. These tasks can be defined with varied degrees of granularity and abstraction [12, 56]. Finally, by *well-justified* we mean that sufficient evidence exists for the effectiveness and expressiveness [53] of the proposed encodings.

In the following sections we demonstrate how we designed our diagrams in the context of this theory of *minimal* and *effective* SQL visualizations. When possible we illustrate our design with figures. We refer the interested reader to our web page queryvis.com with an online demo of QUERYVIS.

4.2 Visual diagrams

We now formally define our visual diagrammatic representation, which we refer to as **QUERYVIS**, and the fragment of SQL it supports. We also describe the automatic transformation from SQL to **QUERYVIS**, formalize how to read the diagrams, and discuss some of their desirable properties.

Our diagrams are inspired by a large body of work on **diagrammatic reasoning systems** [22, 41, 67]. Diagrammatic notations themselves are inspired by the influential *existential graph notation* by Charles Sanders Peirce [60] and exploit topological properties, such as enclosure, to represent logical expressions and set-theoretic relationships. All these representation systems share an origin in that they portray *fragments of first-order-logic* (FOL) through their visual encoding, but at varying degrees of logical expressiveness

and *mainly for monadic relations*. Since SQL (without grouping and null values) essentially represents a tractable FOL fragment between tables and their attributes, QUERYVIS focuses on capturing the FOL representation of an SQL query and representing its underlying logic visually for efficient and intuitive interpretation. We incorporate established visual metaphors from diagrammatic reasoning *into relational schemas* and adapt them where necessary or appropriate.

4.3 Visualizing Conjunctive Queries

We first motivate the design of our diagrams for conjunctive queries, e.g., the one in Fig. 3a and its visualization in Fig. 2a, then generalize to nested queries.

4.3.1 Visualization Design & Effectiveness. Based on the user tasks and abstract properties we identified in Sections 3.1 and 3.2 for conjunctive queries, we reason over the possible mark and channel choices and select the most effective one(s) for each property. Our choices are *effective* as justified using perception studies and visualization best practices.

(1) Relations. To portray a relation, a line mark (straight or curved) drawn between the associated attributes (several shown in Fig. 2a) can simply and clearly encode their connection [57]. In the *Join predicates* task we must distinguish between different comparison operators, which can be done efficiently by placing a label on top of a line to display the applied operator. E.g., the $<>$ operator is shown in Fig. 1 (b) between the L1 and L2 tables. Instead of labels, one could use different line styles (e.g., dashed, double, thin). However, since there are six different operators, this would impose significant learning overhead for a new user and would be less intuitive than labels. To further minimize our design, since the most common type of join is an equijoin, we omit the = label for lines representing equijoins, i.e., unlabeled lines denote an equijoin. Moreover, since the order of elements matters for some operators such as $\{<, \leq, \geq, >\}$, we add an arrowhead mark when necessary to indicate the correct reading order (not illustrated).

For the *Selection predicates* task a line is not an effective encoding as the relation is within one element and needs no portrayal of a connection with another. A constant qualification is better portrayed *in place*, stated explicitly in a row of the referencing table which is highlighted to indicate the presence of a qualification (not illustrated).

(2) Grouping. There exists a plethora of ways to visualize groupings of elements [7, 74] that vary based on the amount, type, and relationship of the elements to be grouped. We adopt an explicit encoding of groups to minimize perceptual and spatial ambiguities. For the *Tables & attributes* task we want to distinguish between different tables and identify their relevant attributes. While tables may share some attributes, we deliberately consider each table as a disjoint set

as we want to make the separation between tables clear. To portray this grouping we use an area/bounding box mark which allows us to use the Gestalt principle of enclosure to denote disjoint set membership [47, 57]. Since the attributes are part of a table (i.e., grouped on a per-table basis) the *table composite mark* is a fitting visual abstraction that encapsulates the membership of attributes under a table and provides a nice visual separation between different table objects. The table composite mark is made up by a set of stacked rectangular box marks, as shown in Fig. 2a. The first box/row in the table represents the table name and is filled with a black background and white text (except the SELECT table which uses a lighter background to distinguish it). The remainder of the rows display the relevant attribute names for which there is an associated selection or join predicate.

Notice that the marks and channels we have chosen, while based on our task analysis and abstraction, are in accordance with previous conjunctive query visualizations and relational schemas, i.e., they are “backwards compatible.”

4.3.2 Visualization Minimality. Our diagrams for conjunctive queries are minimal visualizations because the removal of any mark or channel would lead to incomplete visualizations where at least one of the user tasks cannot be achieved unambiguously. Moreover, we aim to maximize the data-to-ink ratio, i.e., the proportion of a graphic’s “ink” devoted to the non-redundant display of data information [72]. We now explore the issue of visualization minimality for our diagrams for conjunctive queries, noting that only 3 marks are used: table composite marks, lines with arrows/labels, and constant-qualification labels.

(1) Table. The table composite mark is fundamental to a diagram as it identifies the tables involved in the query and their relevant attributes. Removing it would make it impossible to interpret joins. Removing the black background of the first table row, which displays the table name, would eliminate the ability to visually distinguish the table name from its attributes. Removing the gray background in the SELECT table would reduce user ability to detect the root of the query. No attribute row may be removed either, because each must either be part of a join or a selection predicate: their removal would make the *Selection predicates* and *Join predicates* tasks impossible. Alternative visual representations could be chosen, but would require the same number of encodings.

(2) Line. An undirected line mark is essential for portraying joins between two attributes involving operators $\{=, \neq\}$ and its removal would make such *Join predicates* tasks infeasible. Other representation alternatives could be textual, such as referencing the joining attribute next to the other attribute on a table. However, this would lose the advantages of a visual representation where a connection can be

“seen” without being carefully “read.” For joins with operators $\{<, \leq, \geq, >\}$, adding an arrow to a directed line mark ensures the identification of operand order and hence is essential. As an alternative one might consider positional encoding, e.g., enforcing a left-to-right and top-to-bottom reading order. However, this is insufficient for more complicated nested queries (e.g., Fig. 1) which will be discussed in the next section. Conversely, adding arrows to each line would erroneously imply directionality.

Adding a label to the line mark for joins other than the equi-join $\{=\}$ ensures identification of the operator. Its removal would make such *Join predicates* tasks infeasible. Alternative solutions such as use of different line styles exist, but would increase learning difficulty.

(3) Constant-qualification labels. Without these labels the *Selection predicates* task would be infeasible. Alternative visual encodings would add clutter and learning difficulty. Removing the highlight color would limit user ability to distinguish a qualification from an attribute.

4.4 Nested SQL queries

Our main focus is to add nesting to conjunctive queries. We say an SQL query is *nested* if it contains at least one subquery. This is a major step as nested queries (in particular, correlated nested queries) can be very hard to interpret and are not even supported by most visual query builders. We focus on the most expressive subqueries, i.e. those inside the WHERE clause and operators EXISTS, NOT EXISTS, IN, NOT IN, ANY or ALL. Figure 4 shows our currently supported SQL fragment. Those queries have the same expressiveness as relational calculus and its set-based interpretation. However, we currently do not support disjunctions and thus refer to this SQL fragment as *nested conjunctive queries with inequalities*. In addition, SQL queries must fulfill two minor restrictions which we define in Section 5 and argue that they are fulfilled by any meaningful *non-degenerate* SQL query.

Notation. A *predicate* has the form exp1 op exp2 where at most one of the exp 's is a constant (e.g., 3 or 'Alice'), and the other(s) are attribute names optionally with table aliases (e.g., table1.attr2). If a predicate has a constant it is a *selection predicate*, otherwise a *join predicate*. Operator op is an element of $\{<, \leq, =, <>, \geq, >\}$. A *query block* consists of SELECT, FROM, and WHERE clauses including therein defined table aliases and predicates. We call the query block at nesting depth 0 the *root query block*. The *scope* of a query block is the set of query blocks for which the table aliases defined within it are valid (e.g., as shown by the brackets on the right of Fig. 1a). The *root of the scope* is the query block itself. Fig. 1a indicates those by brackets on the left along with their respective *nesting depths*. If a subquery has no other

Q::=	SELECT	C [, C, ..., C] *	select clause
	FROM	S [, S, ..., S]	from clause
	[WHERE	P]	where clause
C::=	[T.]A		column or attribute
S::=	T [AS T]		table (table alias)
P::=	P [AND P ... AND P]		conjunction of predicates
	C O C		join predicate
	C O V		selection predicate
	[NOT] EXISTS (Q)		existential subquery
	C [NOT] IN (Q)		membership subquery
	C O {ALL ANY} (Q)		quantified subquery
O::=	< ≤ = <> ≥ >		comparison operator
T::=			table identifier
A::=			attribute identifier
V::=			string or number

Figure 4: Grammar of supported SQL fragment. Statements enclosed in [] are optional; statements separated by | indicate a choice between alternatives.

nested subquery then its scope and root are the same (e.g., subqueries involving table names L4 and L6 in Fig. 1a).

4.5 Visualizing Nested SQL queries

In order to visualize nested conjunctive queries with inequalities, we must extend the visualizations for conjunctive queries (Section 4.3.1) to further enable the *Quantifiers* and *Nesting order* tasks (Section 3.1). To this end, we are extending the conjunctive-query visualization to allow up to depth 3 nested queries. We illustrate the design for nested queries using Fig. 3b and its associated visualizations Figs. 2b and 2c.

4.5.1 Visualization Design & Effectiveness. We extend the design of our diagrams by keeping the same encodings for the tasks examined in Section 4.3.1 and choose the most effective additional marks and channels for the *Quantifiers* and *Nesting order* tasks from the grouping and hierarchy abstractions, respectively.

(1) Grouping. For the *quantifiers* we must operate a level above the *Tables & attributes* task to group tables based on a quantifier. However, we can leverage the same principles [7, 47, 57, 74] to design an area/bounding box mark that encloses a set of tables. To be distinct from table boxes we use a rounded rectangle mark. As we only have two quantifiers to encode $\{\exists, \forall\}$, we choose to use a dashed and double line style, respectively, so as to avoid labeling. \exists dashed lines are shown in Fig. 2b while a simplified representation with \forall double lines is shown in Fig. 2c.

(2) Hierarchy. A hierarchy can be effectively visualized as a rooted tree or similar node-link structure [57, 71]. For the *Nesting Order* task we could portray the nesting with a logic tree like we introduce in Fig. 5. This would necessitate two arrow types: one to represent the nesting order of subqueries as in logic trees and another to represent the table joins as we do for conjunctive queries (Section 4.3.1). Recall from Section 4.3.2 that positional encoding alone (also shown by Fig. 5 for logic trees) is insufficient for portraying more complicated nested queries (e.g., Fig. 1). Another approach

would be to visually nest sets of tables in bounding boxes, but this would likely lead to cluttered visualizations.

Below we will show that simply by using arrow rules we provide for the reading order (Section 4.6), we can always recover the correct nesting order of each table in an SQL query (Section 5). Hence *additional marks encoding nesting would be redundant* as long as we ensure that arrows are appropriately added to the line marks for *Join predicates* to implicitly encode nesting order. Specifically, we (a) must use directed edges (lines with arrows) for equijoins of tables that are at different nesting depths and (b) determine the direction of the arrow solely by the arrow rules and not the order of attributes around an operator. As an example of the latter constraint, assume we have a join condition $A.attr1 > B.attr2$ where a directed line drawn from A to B denotes operator order. However, table B is a parent of table A in the nesting and thus the arrow rules state the directed edge must be drawn $B \rightarrow A$. Thus we must rewrite the join with the equivalent condition $B.attr2 < A.attr1$.

4.5.2 Visualization Minimality. Just as we did for conjunctive queries in Section 4.3.2, we show that our visualization is minimal also for *nested conjunctive queries with inequalities*. The removal of any of its marks or channels makes it infeasible to unambiguously achieve all the user tasks. In particular, we discuss the two additional and modified marks: lines with arrows/labels and bounding boxes.

(1) Bounding Box. A rounded rectangle mark encloses a query block, i.e. all tables to which a quantifier is applied. Removing it would make the *Quantifiers* task infeasible. One alternative to using enclosure would be to attach the quantifier and a block label or color to each table, but this would require the addition of more marks/channels and would not provide the user with the *at-a-glance* grouping of enclosure.

(2) Line arrows/labels. Besides identifying the pair of attributes involved in a join to support the *Tables & attributes* task, the lines and their arrows/labels are also required for the *Nesting order* task. Removing the lines or their component arrows/labels would make these tasks infeasible. However, the arrow rules (Section 4.6) can be used to unambiguously determine the nesting order of the subqueries in the SQL query with this limited additional encoding as we show in Section 5.2. Alternative approaches, such as positional encoding or nested enclosure, would be insufficient in many cases, and would likely require many more marks and thus more “ink” [72] to display the same data.

4.6 Reading Order

QUERYVIS diagrams are read by starting from the SELECT table and following a depth-first traversal with restarts from unvisited source nodes (i.e. those without incoming arrows).

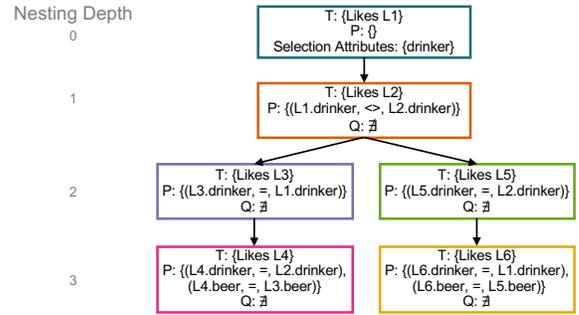


Figure 5: LT representation of the SQL query from Fig. 1a.

Assume an edge goes from $S.attr1$ to $T.attr2$, there is a quantifier $\#$ applied to T and the edge is labeled with comparison operator $<$. Then we can interpret that as follows: *Find attr1 from S s.t. there does not exist any tuple in T where $S.attr1 < T.attr2$.* Recall that unlabeled edges represent equijoins. Also notice that if the two tables are from the same query block, then they are treated as if T has the \exists quantifier applied. Once we finish interpreting an edge, we need to add an AND in our interpretation because we represent a conjunction of predicates.

4.7 Transforming SQL into diagrams

Here we provide an overview of the diagram creation process. We first convert an SQL query into tuple relational calculus (TRC), which is a well-studied transformation to FOL [21]. As a consequence of transforming into FOL, we consider set semantics, 2-valued logic (no NULLs), and no aggregate functions. Moreover, we connect multiple predicates only by conjunctions (i.e., no disjunctions are allowed). In FOL we *no longer have to deal with the various syntactic variants of SQL operators* which do not add expressiveness. This means that operators such as IN, NOT IN, or ALL would be converted to the corresponding FOL quantifiers \exists , $\#$ or \forall .

Logic Tree (LT). Instead of using the TRC representation of a query, it becomes easier to reason over an equivalent representation that makes the *nested scopes of the quantifiers* explicit in the form of a tree, which we denote as *Logic Tree* (LT). It is a rooted tree with each node representing a query block. The *root node* represents the root query block, and the tree structure encodes the nesting hierarchy, i.e., tables and attributes of a node can be referenced in any subtree. Each node in the LT holds the following information:

- (1) *Tables (T)*: The set of tables (or table aliases) defined in its root of the scope;
- (2) *Predicates (P)*: The set of predicates used in the query block. Multiple predicates in the set are related by a conjunction (i.e., $predicate1 \wedge predicate2$ etc.);
- (3) *Quantifier (Q)*: The quantifier applied to the predicates including its negation. This is either \exists , $\#$ or \forall .

Moreover, for the root node of the LT, we also specify the attributes in its select-list (see Fig. 5).

Logic Simplifications. In SQL, queries with universal quantifiers (e.g., our example from Section 1.1) are expressed through nested NOT EXISTS subqueries, which makes them difficult to read. We simplify a LT with nested \nexists quantifiers by applying a standard logical transformation to turn them into \forall quantifiers. In particular, if a LT node ψ has \nexists as its quantifier and only has one child node ψ' that also has \nexists as its quantifier then we can transform ψ to have a \forall quantifier and ψ' to have an \exists quantifier.

To see why, consider following transformation where we apply De Morgan’s law $\neg\exists x \in X.(P(x)) \equiv \forall x \in X.(\neg P(x))$ with $P(x)$ being a propositional logic expression and $\neg a \vee b \equiv a \rightarrow b$ [32] on two LT nodes. Here we write T and S for a set of tables in two query blocks:

$$\neg\exists S.(p_1 \wedge \dots \wedge p_k \wedge \neg\exists T.(p_{k+1} \wedge \dots \wedge p_{k+\ell})) \quad (1)$$

$$\forall S.\neg((p_1 \wedge \dots \wedge p_k) \wedge \neg\exists T.(p_{k+1} \wedge \dots \wedge p_{k+\ell})) \quad (2)$$

$$\forall S.((p_1 \wedge \dots \wedge p_k) \rightarrow \exists T.(p_{k+1} \wedge \dots \wedge p_{k+\ell})) \quad (3)$$

As example, applying this transformation to the LT shown in Fig. 5, led to our diagram from Fig. 1b.

Our Diagram (QUERYVIS). We then transform an LT into our diagram. Starting from the root node we traverse the LT tree in a breadth-first manner and we process each node as follows: (1) Create a table with its attributes for each table defined in the node. (2) Create an appropriate bounding box over the created tables based on the quantifier applied to the node. (3) Update the tables to express the selection predicates by directly writing them into a new row in the table. (4) Create arrows (directed edges) between attributes of tables for which there is a join predicate.

The interesting part is *determining the arrow direction* in step (4): If the join predicate is between two tables T_1 and T_2 from two different LT nodes ψ_1 and ψ_2 , respectively, then we resolve the arrow direction as follows:

- (1) If node ψ_1 is a parent of node ψ_2 then the arrow points from an attribute in T_1 to an attribute in T_2 ;
- (2) Else the arrow points from an attribute in T_2 to an attribute in T_1

Finally, label the arrow based on the join operator op unless op is “=”. As we show in Section 5.2 and Section 3, this seemingly arbitrary choice of arrow directions allows us to create diagrams that (i) are *unambiguous w.r.t. their logic meaning*, (ii) are *minimally verbose*, and (iii) have a *natural reading order* implied by the arrows (cp. to Fig. 1b).

4.8 Diagram Properties

Our diagrams are designed to enhance understanding of an existing SQL query, not to create a query from scratch or

to debug it. Hence they provide the user with simple visual abstractions that ensure the following properties:

(1) Existing metaphors as starting point. Most database users have seen relational schema diagrams before. A simple conjunctive query should not be visualized much differently from a database schema representation. We started from UML and its familiar elements for data modeling and then added visual elements only as needed. *Conjunctive queries* such as Fig. 3a have the lowest visual complexity and are represented as shown in Fig. 2a. In contrast to the SQL representation of the query, in our diagrams *no aliases are needed*. Notice how the lines between the different attributes visualize the join operation and its conditions.

(2) Minimal visual complexity. QUERYVIS uses a minimal number of visual elements (text, line, arrow, table and bounding box) as we describe in greater detail in Section 3. Since our goal is query interpretation, not query specification, we are able to abstract away details such as NULL value handling, which do not impact the intent of the query.

(3) First-order logic representation. QUERYVIS adapts visual metaphors from diagrammatic reasoning into relational schemas. Fig. 3b shows a more complex nested query; its diagram is shown in Fig. 2b. The natural language translations of this query and of the conjunctive query in Fig. 3a have approximately the same length. Similarly, Fig. 2b shows an only slightly increased visual complexity (13% more visual elements) over Fig. 2a. In contrast, the SQL text is much more complex (167% more words), which is a known deficiency of SQL that was aptly termed *poor syntactic locality* [2] and pointed out earlier [35]. We further simplified the visual representation by adding a universal quantifier (\forall), a construct that does not exist in SQL. For example, the representation from Fig. 2b can be further simplified to the one in Fig. 2c by applying the transformation described in Section 4.7. This representation now has only 7% more visible elements than the conjunctive query from Fig. 2a.

(4) Reading-order. Another concept borrowed from diagrammatic reasoning is a *default reading order* [33]. Note from Fig. 2c how the arrows between the relations correspond to the natural language translation. Without arrows, there would be no natural order placed on the existential and universal quantifiers.

5 UNAMBIGUOUS DIAGRAMS

We show that QUERYVIS diagrams for practically relevant SQL queries are provably *unambiguous*, i.e., every diagram can be mapped back to a unique LT. This property is crucial as otherwise a diagram could allow different interpretations.

By construction, each LT node’s content (table aliases, predicates, quantifier) can be directly inferred from the corresponding tables and their connections and bounding boxes

in the diagram. Hence we only have to show that all parent-child relationships between LT nodes can also be recovered. This is not obvious, because we decided to forego an explicit encoding of the LT hierarchy in the QUERYVIS diagram to keep the visual structure simple and to make our diagrams easier to read. Not surprisingly, our supported SQL fragment is powerful enough to create queries that have structurally different LTs, but which map to the same diagram. On the other hand, for each meaningful SQL query we observed in practice it is possible to map its QUERYVIS diagram back to a unique LT. We will formalize the common properties of these practical queries, which we call *non-degenerate*.

5.1 Non-degenerate SQL Statements

Consider again the unique-set query in Fig. 1a and observe the following properties. First, each join condition references a relation whose alias appears in the FROM clause of the same query block. For instance, L1.drinker <> L2.drinker appears in the block introducing L2, even though it could in theory also appear at a deeper nesting depth. The latter does not make sense in practice, as the user will attempt to define a condition “where it belongs,” i.e., as early as possible. Second, note how each table alias introduced at shallower depth is referenced by a join condition at a deeper level. This again makes sense, because the main purpose of introducing a complex nesting structure is to relate attributes from an outer block with attributes in nearby inner blocks.

We formalize these commonly observed properties below and then show that queries with those properties result in unambiguous QUERYVIS diagrams.

PROPERTY 5.1 (LOCAL ATTRIBUTES). *Each predicate in a query block references at least one local attribute, i.e. an attribute of a table from the same query block.*

Property 5.1 implies that the predicate is placed as high as possible in the nesting hierarchy of the query and it is not possible to pull it up to an ancestor in the LT. If this property does not hold, then the predicate violating this property actually expresses a *disjunction*. For example, consider the following query written in TRC:

$$\{F.person \mid F \in Frequent \wedge \neg \exists S \in Serves.(S.bar = F.bar \wedge F.bar = \text{“Owl”})\}$$

This query is in the SQL fragment of Section 4.4 but violates **Property 5.1**: the selection predicate $F.bar = \text{“Owl”}$ could be pulled up to nesting depth 0, and after applying De Morgan’s law on the expression we get a disjunction:

$$\{F.person \mid F \in Frequent \wedge (F.bar \neq \text{“Owl”} \vee \neg \exists S \in Serves.(S.bar = F.bar))\}$$

PROPERTY 5.2 (CONNECTED SUBQUERIES). *Each nested query block q_i either has a predicate referencing an attribute from its parent query block, or each of its directly nested query blocks references both q_i and its parent.*

Property 5.2 ensures that there are meaningful logical connections between a query block and its nested subqueries.

5.2 Proof of unambiguity

In addition to being non-degenerate, the queries we observe in practice also do not have more than 3 levels of nesting. Hence we call a diagram *valid* if there exists a non-degenerate SQL query up to nesting depth 3 that is part of the SQL fragment discussed in Section 4.4 that maps to it. We next show that any valid diagram can be uniquely interpreted.

PROPOSITION 5.1 (UNAMBIGUITY). *For any valid QUERYVIS diagram there exists exactly one LT that maps to it.*

We confirm **Proposition 5.1** by considering all possible valid diagrams up to nesting depth 3. Given the arrow rules in Section 4.6 one can uniquely identify the corresponding LT node for any given query block in our diagram.

6 EXPERIMENTAL EVALUATION

We designed a user study to test whether our diagrams help users understand SQL queries *in less time* and *with fewer errors*, on average. We thus tested two **hypotheses**:

(H1): Study participants can understand queries *in less time* with our diagrams than by reading SQL code alone.

(H2): Study participants can understand queries *with fewer errors* with our diagrams than by reading SQL code alone.

The study design and analysis plan was preregistered before we conducted the experiment and is available on OSF at osf.io/mycr2 Details on the study are also available in the full version on this paper [51].

6.1 Study Design

To test our hypotheses, we designed an easily-scalable *within-subjects study* (i.e., all study participants were exposed to all query interfaces [64]) and made it available for 3 weeks from Jan 24, 2020–Feb 13, 2020 on Amazon Mechanical Turk (AMT). During that time, we recruited $n = 42$ legitimate participants. In the study, we presented 9 queries to participants in one of 3 conditions: (1) seeing a query as SQL alone (*SQL*), (2) seeing a query as a logical diagram that was generated from SQL (*QV*), or (3) seeing both SQL and QV at the same time (*Both*). We then tracked the time needed and errors made by each participant while trying to find the correct interpretation for each query. Note that the entire study included 12 queries yet 3 of them are related to an extension of our visualization under development (Groub-By queries) and are not presented here. The results of analyzing all 12

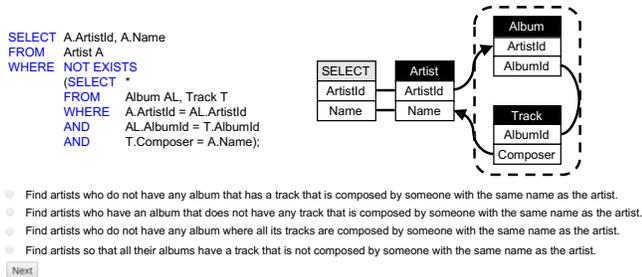


Figure 6: Example query from our study. The query is shown in the *Both* condition, in which a participant sees the query in both SQL (left) and our QV diagram (right).

questions are similar to analyzing only the 9. Please see our supplemental material at osf.io/mycr2 for details.

Multiple-choice questions. Our study consisted of 9 multiple-choice questions (MCQs) Q_1 – Q_9 . Each MCQ asked the participant to choose the best interpretation for a presented query from four choices. Following best practices in MCQ creation [80], all 4 of the choices were designed to read very similar to each other so that a participant with little knowledge of SQL would be incapable of eliminating any of the 4 choices. Upon answering a question we would provide immediate feedback to the participant by highlighting the correct answer. All 9 of our questions were based on the widely used Chinook database schema [20]. Our 9 questions were split into 3 categories: conjunctive with no self-joins, conjunctive with self-joins and nested queries where each category had 3 questions (3 questions fit well in our Latin square design of 3 conditions). We chose those 3 categories based on past work [4, 11] that analyzes the most prevalent types of SQL errors. Each category consisted of 3 queries: one simple, one medium and one complex. The complexity was designated based on the number of joins and number of table aliases referenced in the query. Figure 6 shows the interface for the condition *Both* for one of the 9 questions.

Latin square design. For our within-subjects study we adopted a *Latin square design* [50, 55] by which each participant experienced all three conditions (SQL, QV, and Both) in a particular randomized way that reduces potential biases in our analysis due to condition ordering effects. Each participant answered *all 9 questions in the same order* but their condition for each question was different and depended on their *sequence number* in $\{S_1, S_2, \dots, S_6\}$. The sequence number is the order in which conditions appear and is based on a repeating triplet of conditions. There are in total 6 sequences, one for each of the 6 possible permutations in a triplet (e.g., sequence S_1 is “SQL \rightarrow QV \rightarrow Both”, and S_2 is “SQL \rightarrow Both \rightarrow QV”). Once a participant completes the first triplet of questions the permutation repeats. As a result, *each participant experiences each condition in 3 different questions* (e.g., a participant in sequence #1 sees questions $\{Q_1, Q_4, Q_7\}$

presented as SQL). We assigned a sequence number to each participant in a round robin fashion and ensured a balanced number of participants in each sequence.

Recruitment of study participants. Upon approval from our Institutional Review Board, we recruited participants from *Amazon Mechanical Turk* (AMT). Informed consent was gained from all participants before they accepted the task. We restricted participants to be from the USA, since those workers have literacy in English (which is the language our test was written in), and the USA is the country with the largest number of AMT workers [28]. We also restricted participants to those with a 95% *approval rating* on AMT, which is a typical requirement in order to reduce potential speeders (i.e., Turk workers that speed through tasks in the hopes of fulfilling the minimum requirements). All our study participants had to pass an *SQL qualification exam* to ensure that they had at least a basic proficiency with SQL. The qualification exam consisted of 6 SQL questions, took at most 10 min, and workers needed at least 4/6 correct answers.

We conducted a $n = 12$ pilot study followed by the full study with $n = 42$ AMT workers. An additional 38 *illegitimate* workers were excluded from the analysis because of obvious speeding or cheating behavior (see the supplemental material for details). The 42 *legitimate* participants included 15 women, 24 men, and 3 that didn’t disclose their gender. Participants came from a variety of different professions ranging from software engineers and college students to homemakers. Each worker could only take the test once and we prevented workers from taking the test if they had participated in our pilot studies.

Visual diagram tutorial. Once workers passed the qualification exam and accepted our task, they were given a self-paced six-page tutorial about our logical diagrams. The tutorial introduced our basic visual notations by showing SQL examples and their diagrams. The mean (resp. median) time spent on the tutorial was approx. 3 (resp. 2) minutes. Note that our participants potentially had many years of experience with SQL (In a post-test survey, the median SQL usage frequency was 2 times per month), but *only a few minutes of experience with our logical diagrams*. The tutorial is available in our supplemental material.

Performance-based monetary incentivisation. Once they finished the tutorial, participants were required to answer at least 5 test questions correctly within 50 minutes in order for their work to be accepted and for them to earn a base pay of \$5.20 USD. The time limit was placed in order to encourage participants to take the test in one sitting. This is critical because one of our performance metrics is the *participant’s time spent per question*, and it is not uncommon for AMT workers to accept multiple tasks with long time limits at the same time and do them in parallel [16]. The base pay was determined following standard guidelines [76] based on

mean pilot test duration and a USA living wage of \$15/hr. To incentivize the workers to complete the test *efficiently and correctly* we used a staggered monetary incentivization scheme that added additional bonuses for having more questions correct in shorter time. Please see the online materials on OSF at osf.io/mycr2 for details.

6.2 Study Analysis

We conducted a power analysis with timing data from a pilot study to estimate the required number of participants. We then used traditional null hypothesis significance testing with corrections for multi-hypothesis testing [9] as well as interval estimation of effect sizes [24, 29].

Required sample size through power analysis. In order to test *the effect of each condition on time*, we ran a pilot and conducted a power analysis [55, 79] to *estimate the number of participants we needed*. We did not do power analysis on error due to the discrete nature of the data. The power of a statistical test is the pre-study probability of correctly rejecting the null hypothesis, e.g., detecting a difference between conditions when one exists [38]. The rejection threshold associated with the p -value is denoted α and the power is denoted $1 - \beta$. Our power analysis assumes comparing two-sample means with a one-tailed test given parameters of $\alpha = 5\%$ and $1 - \beta = 90\%$. Our two-sample means were the measured *mean times* for conditions *SQL* and *QV* based on the $n = 12$ pilot study. We conducted a one-tailed test with the alternative hypothesis that the mean time of the *QV* condition is less than the mean time of the *SQL* condition, in contrast to the null hypothesis which was the reverse statement. We chose not to use mean error results in our power analysis as the values are discretized due to the small number of questions per condition and thus unsuitable for a one-tailed test. Based on the pilot, the estimated sample size required to achieve the desired power was $n = 84$, rounded up to the nearest multiple of six to ensure an even split of participants across sequences. During the 3 weeks the study was available on AMT, we were only able to recruit 42 legitimate participants for which we report results below.

Statistical tests for rejecting the null. We calculated 6 quantities: The mean time per question in seconds and the mean error rate in $[0, 1]$, across all three conditions. We had 2 hypotheses for time and 2 hypotheses for error:

Time: (1) $\text{time}_{QV} < \text{time}_{SQL}$, and (2) $\text{time}_{Both} < \text{time}_{SQL}$.

Error: (1) $\text{err}_{QV} < \text{err}_{SQL}$, and (2) $\text{err}_{Both} < \text{err}_{SQL}$.

As we conducted two hypothesis tests on the same data (i.e., two hypotheses on time, and separately two on error), we *adjusted all p -values using the Benjamini and Hochberg procedure* [9] in order to minimize false discoveries caused by multiple hypothesis testing. We examined the distributions of the collected data to ensure that the assumptions of our

statistical tests are not violated. We tested the normality of the distribution for each condition using a Q-Q plot and Shapiro-Wilk test [65] with $\alpha = 5\%$. We determined that the data was not normally distributed, as is required for a *parametric statistical test* [66]. Moreover, the distributions were not all transformable to normal using the same exponent via a Box-Cox transformation [63]. Thus, we used *non-parametric statistical tests* [66] which do not require our data to be normally distributed. We separately ran one-tailed Wilcoxon signed-rank [77] tests for our four hypotheses.

Notice that, as a within-subjects-study, we are comparing how the performance (time and error rate) of a participant varied between the 3 conditions. This analysis is done *for each participant individually* before averaging across participants.

Interval estimation of effect size. Estimation involves reporting effect sizes with interval estimates, enabling more nuanced interpretations [24, 29]. We augment our analyses with estimation because traditional dichotomous inference [10] or dichotomania [37] in null hypothesis significance testing — e.g., using only $p < 0.05$ as a cutoff for “significant” results — can lead to erroneous beliefs and poor decision making [10, 37, 38, 75]. A large p -value does not imply evidence in favor of the null hypothesis. Instead, a p -value should be interpreted as how compatible the null hypothesis is with the data and any $p < 1$ indicating some incompatibility [38, 75].

In particular, we report the sample median completion time and sample mean error and use the difference between the sample means/medians as the effect size [23, 30]. We use bias-corrected and accelerated (BCa) 95% *confidence intervals (CIs)* to indicate the *range of plausible values* for the mean time and mean error [29, 31].

6.3 Study Results and Discussion

The results of our study on the 42 legitimate participants are summarized in Fig. 7 and the red boxes below.

Time.

- (1) *There is strong evidence that participants are meaningfully faster (-20%) using QV than SQL ($p < 0.001$)*
- (2) *Participants take a similar amount of time (-1%) using Both and SQL ($p = 0.30$)*

Error rate.

- (1) *There is weak evidence that participants make meaningfully fewer errors (-21%) using QV than SQL ($p = 0.15$)*
- (2) *There is weak evidence that participants make meaningfully fewer errors (-17%) using Both than SQL ($p = 0.16$)*

Note that the p -values above are adjusted to account for multiple hypothesis testing.

Validation, speeders, cheaters. In our study design and preregistration, we foresaw the problem of speeders. However, we discovered that some users cheated. Near the end

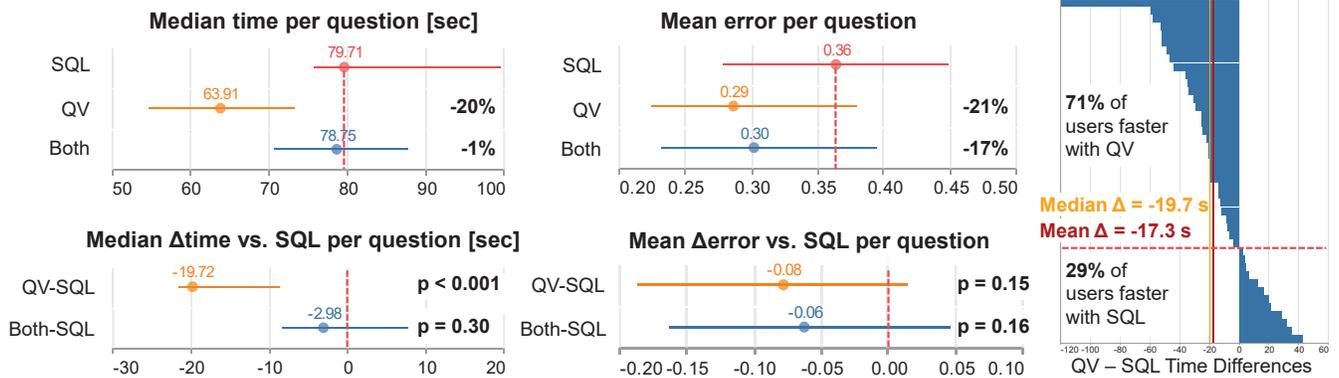


Figure 7: Median time per question and mean error rates for each of the 3 conditions with their respective 95% BCa [31] confidence intervals (CIs). Participants were on average faster using QV than SQL (-20%, $p < 0.001$), and slightly faster using Both than SQL (-1%, $p = 0.30$). Participants made less errors using QV than SQL (-21%, $p = 0.15$) and less errors with Both than SQL (-17%, $p = 0.16$). Note that the p -values are based on within-subjects statistical tests, thus confidence intervals of the mean could overlap substantially without indicating a high p -value [29]. The p -value for a given hypothesis tests an individual’s performance in the QV or Both condition vs. their performance in the SQL condition. The p -values are listed along with the respective percentage difference of the condition’s values vs. their counterpart SQL value. However, the confidence intervals capture the median/mean of a given condition for all participants – neglecting analysis at a within-participant level. The bottom row shows the per-participant differences of QV or Both condition vs. their performance in the SQL condition.

of the study, many participants had all questions correct but needed less than 30 seconds to solve them. We excluded 38 users we believe cheated in our study (see supplemental material for details).

Future potential from regular use. Our user study provides strong evidence that our novel logical diagrams helped participants interpret queries more quickly than using SQL alone, and good evidence that combining our diagrams together with SQL helped them interpret queries with fewer errors. We would like to emphasize that the study participants had only a few minutes of exposure to our logical diagrams, but potentially many years of experience with SQL. Moreover, from a post-test survey about their frequency of SQL use, the median reported frequency was a couple times per month. We may further reduce user time and error rate by providing an interactive rather than static tutorial [49].

Logical SQL patterns. Since the notion of *patterns* was formulated by the architect Christopher Alexander in the 1970s [6], *patterns*, *archetypes*, *frameworks*, *templates* have been successfully used in various sciences, e.g., the widely-used *design patterns* in software engineering [34]. Despite having taught SQL to hundreds of student, we are not aware of any current pedagogic approach for systematically teaching SQL using *logical SQL design patterns*. We are also not aware of any ongoing efforts to apply the rich literature on logical diagrams to the task of abstracting SQL queries using *logical SQL patterns*, which may help users quickly understand the semantics and meaning of queries. We speculate that for understanding query intent, *the potential to further improve speed and reduce error rates through learning and*

regularly using logical SQL patterns is enormous. We hope that further studies can verify this hypothesis. We also hope that our preliminary results can serve as a call to action for the database and visualization communities to investigate novel design and education approaches for *understanding SQL queries and their patterns*.

7 CONCLUSIONS

SQL queries can be challenging to read and understand, even for experts. Code maintenance, query modification, and query reuse require users to *read existing SQL queries*, both their own and those written by others. To facilitate this task, we have described an approach that can automatically transform a large fragment of SQL queries into *unambiguous*, and *visually minimal* visual diagrams. Those diagrams are based on a first order logic interpretation of SQL and thus *reflect and expose the logical patterns* innate to SQL. Our diagrams significantly improve the speed of understanding SQL queries even with virtually no prior exposure to our diagrams by participants. In future work we will investigate visual metaphors and extend the supported fragment of SQL, in particular for groupings and arithmetic predicates.

ACKNOWLEDGMENTS

This work was supported in part by a Khoury seed grant program, the National Science Foundation (NSF) under award numbers CAREER IIS-1762268 and ACI-1640575. We also like to thank Jonathan Danaparamita for singlehandedly implementing the original QUERYVIS prototype.

REFERENCES

- [1] Azza Abouzied, Joseph M. Hellerstein, and Avi Silberschatz. 2012. DataPlay: interactive tweaking and example-driven correction of graphical database queries. In *UIST*. ACM, 207–218. <https://doi.org/10.1145/2380116.2380144>
- [2] Azza Abouzied, Joseph M. Hellerstein, and Avi Silberschatz. 2012. Playful Query Specification with DataPlay. *PVLDB* 5, 12 (2012), 1938–1941. <https://doi.org/10.14778/2367502.2367542>
- [3] Active Query Builder. 2019. <https://www.activequerybuilder.com/>.
- [4] Alireza Ahadi, Vahid Behbood, Arto Vihavainen, Julia Prior, and Raymond Lister. 2016. Students’ Syntactic Mistakes in Writing Seven Different Types of SQL Queries and Its Application to Predicting Students’ Success. In *Technical Symposium on Computing Science Education (SIGCSE)*. ACM, 401–406. <https://doi.org/10.1145/2839509.2844640>
- [5] Javad Akbarnejad, Gloria Chatzopoulou, Magdalini Eirinaki, Suju Koshy, Sarika Mittal, Duc On, Neoklis Polyzotis, and Jothi S. Vin-dhiya Varman. 2010. SQL QueRIE Recommendations. *PVLDB* 3, 1 (2010), 1597–1600. <https://doi.org/10.1145/2839509.2844640>
- [6] Christopher Alexander, Sara Ishikawa, and Murray Silverstein. 1977. *A pattern language: towns, buildings, construction*. Oxford University Press, New York. <http://www.loc.gov/catdir/enhancements/fy0603/74022874-d.html>
- [7] Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter Rodgers. 2014. Visualizing Sets and Set-typed Data: State-of-the-Art and Future Challenges. In *EuroVis - STARs*. The Eurographics Association. <https://doi.org/10.2312/eurovisstar.20141170>
- [8] Leilani Battle, Danyel Fisher, Robert DeLine, Mike Barnett, Badrish Chandramouli, and Jonathan Goldstein. 2016. Making Sense of Temporal Queries with Interactive Visualization. In *CHI*. ACM, 5433–5443. <https://doi.org/10.1145/2858036.2858408>
- [9] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- [10] Lonni Besançon and Pierre Dragicevic. 2019. The Continued Prevalence of Dichotomous Inferences at CHI. In *CHI Extended Abstracts (CHI EA)*. ACM, 1–11. <https://doi.org/10.1145/3290607.3310432>
- [11] Stefan Brass and Christian Goldberg. 2006. Semantic errors in SQL queries: A quite complete list. *Journal of Systems and Software* 79, 5 (2006), 630 – 644. <https://doi.org/10.1016/j.jss.2005.06.028> Quality Software.
- [12] M. Brehmer and T. Munzner. 2013. A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2376–2385. <https://doi.org/10.1109/TVCG.2013.124>
- [13] Tiziana Catarci. 2000. What Happened When Database Researchers Met Usability. *Inf. Syst.* 25, 3 (2000), 177–212. [https://doi.org/10.1016/S0306-4379\(00\)00015-6](https://doi.org/10.1016/S0306-4379(00)00015-6)
- [14] Tiziana Catarci, Maria Francesca Costabile, Stefano Levialdi, and Carlo Batini. 1997. Visual Query Systems for Databases: A Survey. *J. Vis. Lang. Comput.* 8, 2 (1997), 215–260. <https://doi.org/10.1006/jvlc.1997.0037>
- [15] Claudio Cerullo and Marco Porta. 2007. A System for Database Visual Querying and Query Visualization: Complementing Text and Graphics to Increase Expressiveness. In *International Workshop on Database and Expert Systems Applications (DEXA)*. IEEE, 109–113. <https://doi.org/10.1109/DEXA.2007.91>
- [16] Jesse Chandler, Pam Mueller, and Gabriele Paolacci. 2014. Nonnaïveté among Amazon Mechanical Turk workers: Consequences and solutions for behavioral researchers. *Behavior research methods* 46, 1 (2014), 112–130. <https://doi.org/10.3758/s13428-013-0365-7>
- [17] Bikash Chandra, Bhupesh Chawda, Biplab Kar, K. V. Reddy, Shetal Shah, and S. Sudarshan. 2015. Data Generation for Testing and Grading SQL Queries. *The VLDB Journal* 24, 6 (Dec. 2015), 731–755. <https://doi.org/10.1007/s00778-015-0395-0>
- [18] Gloria Chatzopoulou, Magdalini Eirinaki, and Neoklis Polyzotis. 2009. Query Recommendations for Interactive Database Exploration. In *SSDBM (LNCS)*, Vol. 5566. Springer, 3–18. https://doi.org/10.1007/978-3-642-02279-1_2
- [19] Chaomei Chen. 2006. *Information visualization: beyond the horizon* (2nd ed.). Springer, New York. <http://www.loc.gov/catdir/enhancements/fy0663/2006920915-d.html>
- [20] Chinook database. 2020. <https://github.com/lerocha/chinook-database>.
- [21] E. F. Codd. 1970. A Relational Model of Data for Large Shared Data Banks. *Commun. ACM* 13, 6 (1970), 377–387. <https://doi.org/10.1145/362384.362685>
- [22] Bruno Courcelle. 1990. Chapter 5 - Graph Rewriting: An Algebraic and Logic Approach. In *Formal Models and Semantics*, Jan Van Leeuwen (Ed.). Elsevier, Amsterdam, 193 – 242. <https://doi.org/10.1016/B978-0-444-88074-1.50010-X>
- [23] Geoff Cumming. 2013. *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge. <https://doi.org/10.4324/9780203807002>
- [24] Geoff Cumming. 2014. The New Statistics: Why and How. *Psychological Science* 25, 1 (2014), 7–29. <https://doi.org/10.1177/0956797613504966>
- [25] Jonathan Danaparamita and Wolfgang Gatterbauer. 2011. QueryViz: Helping Users Understand SQL queries and their patterns. In *EDBT*. ACM, 558–561. [10.1145/1951365.1951440](https://doi.org/10.1145/1951365.1951440)
- [26] dbForge. 2019. <https://www.devart.com/dbforge/mysql/querybuilder/>.
- [27] Dan Diaper and Neville Stanton. 2004. *The Handbook of Task Analysis for Human-Computer Interaction*. Lawrence Erlbaum Associates.
- [28] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and dynamics of mechanical turk workers. In *WSDM*. ACM, 135–143. [10.1145/3159652.3159661](https://doi.org/10.1145/3159652.3159661)
- [29] Pierre Dragicevic. 2016. *Fair Statistical Communication in HCI*. Springer International Publishing, Cham, 291–330. https://doi.org/10.1007/978-3-319-26633-6_13
- [30] Pierre Dragicevic. 2018. Can we call mean differences “effect sizes”? <https://transparentstatistics.org/2018/07/05/meanings-effect-size/>
- [31] Bradley Efron. 1987. Better Bootstrap Confidence Intervals. *J. Amer. Statist. Assoc.* 82, 397 (1987), 171–185. <https://doi.org/10.1080/01621459.1987.10478410>
- [32] Herbert Enderton and Herbert B Enderton. 2001. *A mathematical introduction to logic*. Elsevier. <https://doi.org/10.1016/C2009-0-22107-6>
- [33] Andrew Fish and John Howse. 2004. Towards a Default Reading for Constraint Diagrams. In *Diagrams*. Springer, 51–65. [10.1007/978-3-540-25931-2_8](https://doi.org/10.1007/978-3-540-25931-2_8)
- [34] Erich Gamma. 1995. *Design patterns: elements of reusable object-oriented software*. Addison-Wesley. <https://doi.org/doi/book/10.5555/186897>
- [35] Wolfgang Gatterbauer. 2011. Databases will Visualize Queries too. *PVLDB* 4, 12 (2011), 1498–1501. <http://www.vldb.org/pvldb/vol4/p1498-gatterbauer.pdf> (<http://www.youtube.com/watch?v=kVFnQRGAQls>).
- [36] Sebastian Gehrman, Falcon Z. Dai, Henry Elder, and Alexander M. Rush. 2018. End-to-End Content and Plan Selection for Data-to-Text Generation. In *International Conference on Natural Language Generation (INLG)*. ACM, 46–56. <https://doi.org/10.18653/v1/w18-6505>
- [37] Sander Greenland. 2017. Invited Commentary: The Need for Cognitive Science in Methodology. *American Journal of Epidemiology* 186, 6 (2017), 639–645. <https://doi.org/10.1093/aje/kwx259>
- [38] Sander Greenland, Stephen J. Senn, Kenneth J. Rothman, John B. Carlin, Charles Poole, Steven N. Goodman, and Douglas G. Altman. 2016. Statistical tests, P values, confidence intervals, and power: a guide

- to misinterpretations. *European Journal of Epidemiology* 31, 4 (2016), 337–350. <https://doi.org/10.1007/s10654-016-0149-3>
- [39] Laura M. Haas, Johann Christoph Freytag, Guy M. Lohman, and Hamid Pirahesh. 1989. Extensible Query Processing in Starburst. In *SIGMOD*. 377–388. [10.1145/67544.66962](https://doi.org/10.1145/67544.66962)
- [40] Bill Howe and Garret Cole. 2010. SQL is Dead; Long Live SQL: Lightweight Query Services for Ad Hoc Research Data. In *4th Microsoft eScience Workshop*. <https://homes.cs.washington.edu/~billhowe/projects/2014/03/22/SQLShare.html>
- [41] John Howse. 2008. Diagrammatic Reasoning Systems. In *ICCS (LNCS)*, Vol. 5113. Springer, 1–20. [10.1007/978-3-540-70596-3_1](https://doi.org/10.1007/978-3-540-70596-3_1)
- [42] Yannis E. Ioannidis. 2008. From Databases to Natural Language: The Unusual Direction. In *NLDB (LNCS)*, Vol. 5039. Springer, 12–16. [10.1007/978-3-540-69858-6_3](https://doi.org/10.1007/978-3-540-69858-6_3)
- [43] Hannu Jaakkola and Bernhard Thalheim. 2003. Visual SQL – High-Quality ER-Based Query Treatment. In *ER (Workshops) (LNCS)*. Springer, 129–139. [10.1007/978-3-540-39597-3_13](https://doi.org/10.1007/978-3-540-39597-3_13)
- [44] H. V. Jagadish, Adriane Chapman, Aaron Elkiss, Magesh Jayapandian, Yunyao Li, Arnab Nandi, and Cong Yu. 2007. Making database systems usable. In *SIGMOD*. 13–24. [10.1145/1247480.1247483](https://doi.org/10.1145/1247480.1247483)
- [45] Nodira Khoussainova, Magdalena Balazinska, Wolfgang Gatterbauer, YongChul Kwon, and Dan Suciu. 2009. A Case for A Collaborative Query Management System. In *CIDR*. http://www-db.cs.wisc.edu/cidr/cidr2009/Paper_94.pdf
- [46] Nodira Khoussainova, YongChul Kwon, Magdalena Balazinska, and Dan Suciu. 2010. SnipSuggest: A Context-Aware SQL-Autocomplete System. *PVLDB* 4, 1 (2010), 22–33. <https://doi.org/10.14778/1880172.1880175>
- [47] Kurt Koffka. 2013. *Principles of Gestalt psychology*. Routledge. <https://doi.org/10.4324/9781315009292>
- [48] Georgia Koutrika, Alkis Simitsis, and Yannis E. Ioannidis. 2010. Explaining Structured Queries in Natural Language. In *ICDE*. IEEE, 333–344. [10.1109/ICDE.2010.5447824](https://doi.org/10.1109/ICDE.2010.5447824)
- [49] Bum Chul Kwon and Bongshin Lee. 2016. A Comparative Evaluation on Online Learning Approaches Using Parallel Coordinate Visualization. In *CHI*. 993–997. <https://doi.org/10.1145/2858036.2858101>
- [50] Johannes Ledolter and Arthur J Swersey. 2007. *Testing 1-2-3: experimental design with applications in marketing and service operations*. Stanford Business Books. <https://www.sup.org/books/title/?id=4513>
- [51] Aristotelis Leventidis, Jiahui Zhang, Cody Dunne, Wolfgang Gatterbauer, H.V. Jagadish, and Mirek Riedewald. 2020. QueryVis: Logic-based diagrams help users understand complicated SQL queries faster. *OSF Preprints* (April 2020). <https://doi.org/10.31219/osf.io/btszh> <https://osf.io/btszh>
- [52] Guoliang Li, Ju Fan, Hao Wu, Jiannan Wang, and Jianhua Feng. 2011. DBase: Making Databases User-Friendly and Easily Accessible. In *CIDR*. 45–56. http://cidrdb.org/cidr2011/Papers/CIDR11_Paper6.pdf
- [53] Jock Mackinlay. 1986. Automating the design of graphical presentations of relational information. *TOG: ACM Transactions on Graphics* 5, 2 (April 1986), 110–141. <https://doi.org/10.1145/22949.22950>
- [54] Microsoft Access. 2019. <https://products.office.com/en-us/access>.
- [55] Douglas C Montgomery. 2013. *Design and analysis of experiments* (8th edition ed.). John Wiley & Sons, Inc., Hoboken, NJ. <https://doi.org/10.5555/1206386>
- [56] T. Munzner. 2009. A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (Nov. 2009), 921–928. <https://doi.org/10.1109/TVCG.2009.111>
- [57] T. Munzner. 2014. *Visualization Analysis and Design*. CRC Press. [10.1201/b17511](https://doi.org/10.1201/b17511)
- [58] Brad A. Myers. 1990. Taxonomies of visual programming and program visualization. *Journal of Visual Languages & Computing* 1, 1 (1990), 97 – 123. [https://doi.org/10.1016/S1045-926X\(05\)80036-9](https://doi.org/10.1016/S1045-926X(05)80036-9)
- [59] Christopher Olston, Shubham Chopra, and Utkarsh Srivastava. 2009. Generating example data for dataflow programs. In *SIGMOD*. ACM, 245–256. [10.1145/1559845.1559873](https://doi.org/10.1145/1559845.1559873)
- [60] Charles Sanders Peirce. 1933. *Collected Papers*. Vol. 4. Harvard University Press. <https://doi.org/10.1177/000271623417400185>
- [61] pgAdmin. 2019. <https://www.pgadmin.org/>.
- [62] QueryScope. 2019. <https://sqldep.com/>.
- [63] RM Sakia. 1992. The Box-Cox transformation technique: a review. *Journal of the Royal Statistical Society: Series D (The Statistician)* 41, 2 (1992), 169–178. <https://doi.org/10.2307/2348250>
- [64] Howard J Seltman. 2012. Experimental design and analysis. *Online at: http://www.stat.cmu.edu/~hseltman/309/Book/Book.pdf* (2012).
- [65] Samuel Sanford Shapiro and Martin B Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3/4 (1965), 591–611. <https://doi.org/10.2307/2333709>
- [66] David J Sheskin. 2003. *Handbook of parametric and nonparametric statistical procedures*. Chapman and Hall/CRC. <https://doi.org/10.1201/9781420036268>
- [67] Sun-Joo Shin. 1995. *The Logical Status of Diagrams*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511574696>
- [68] Alkis Simitsis and Yannis E. Ioannidis. 2009. DBMSs Should Talk Back Too. In *CIDR*. http://www-db.cs.wisc.edu/cidr/cidr2009/Paper_119.pdf
- [69] Sloan digital sky survey. 2019. <http://www.sdss.org/>.
- [70] SQL Server Management Studio. 2019. <https://www.microsoft.com/en-us/sql-server/sql-server-downloads>.
- [71] Kozo Sugiyama, Shojiro Tagawa, and Mitsuhiro Toda. 1981. Methods for visual understanding of hierarchical system structures. *IEEE Transactions on Systems, Man, and Cybernetics* 11, 2 (1981), 109–125. <https://doi.org/10.1109/TSMC.1981.4308636>
- [72] Edward R. Tufte. 1986. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, USA. <https://dl.acm.org/doi/book/10.5555/33404>
- [73] Jeffrey D. Ullman. 1988. *Principles of Database and Knowledge-base Systems, Vol. I*. Computer Science Press, Inc., New York, NY, USA. <https://dl.acm.org/doi/book/10.5555/42790>
- [74] Corinna Vehlow, Fabian Beck, and Daniel Weiskopf. 2017. Visualizing Group Structures in Graphs: A Survey. *Computer Graphics Forum* 36, 6 (2017), 201–225. <https://doi.org/10.1111/cgf.12872>
- [75] Ronald L. Wasserstein and Nicole A. Lazar. 2016. The ASA Statement on p-Values: Context, Process, and Purpose. *The American Statistician* 70, 2 (2016), 129–133. <https://doi.org/10.1080/00031305.2016.1154108>
- [76] WeAreDynamo. 2019. Guidelines for Academic Requesters. http://www.wearedynamo.org/Guidelines_for_Academic_Requesters.pdf
- [77] Frank Wilcoxon. 1945. Individual Comparisons by Ranking Methods. *Biometrics Bulletin* 1, 6 (1945), 80–83. <https://doi.org/10.2307/3001968>
- [78] Kun Xu, Lingfei Wu, Zhiguo Wang, Yansong Feng, and Vadim Sheinin. 2018. SQL-to-Text Generation with Graph-to-Sequence Model. In *EMNLP*. ACL, 931–936. <https://doi.org/10.18653/v1/d18-1112>
- [79] Koji Yatani. 2016. *Effect Sizes and Power Analysis in HCI*. Springer International Publishing, Cham, 87–110. https://doi.org/10.1007/978-3-319-26633-6_5
- [80] Dawn M. Zimmaro. 2010. *Writing Good Multiple-Choice Exams*. Center for Teaching and Learning, UT Austin. <https://facultyinnovate.utexas.edu/sites/default/files/writing-good-multiple-choice-exams-fic-120116.pdf>
- [81] Moshé M. Zloof. 1977. Query-by-Example: A Data Base Language. *IBM Systems Journal* 16, 4 (1977), 324–343. <https://doi.org/10.1147/sj.164.0324>