Research online Publishing
www.ronpub.com

Open Access

Open Journal of Semantic Web (OJSW)
Volume 8, Issue 1, 2021

http://www.ronpub.com/ojsw
ISSN 2199-336X

Generating Sound from the Processing in Semantic Web Databases

Sven Groppe, Rico Klinckenberg, Benjamin Warnke

Institute of Information Systems (IFIS), University of Lübeck, Ratzeburger Allee 160, D-23562 Lübeck, Germany
groppe@ifis.uni-luebeck.de, rico.klinckenberg@gmx.net, benjamin.warnke@uni-luebeck.de

Abstract

Databases process a lot of intermediate steps generating many intermediate results during data processing for answering queries. It is not easy to understand these complex tasks and algorithms for students, developers and all those interested in databases. For this purpose, an additional medium is sonification, which maps data to auditory dimensions and offers a new audible experience to their listeners. Hence, we propose a sonification of query processing paired with a corresponding visualization both integrated in a web application. In a demonstration of our approach and in an extensive user evaluation we show that listeners increase their understanding of the operators' functionality and sonification supports easy remembering of requirements like merge joins work on sorted input. Furthermore, new ways of analyzing query processing are possible with our proposed sonification approach.

Type of Paper and Keywords

Regular research paper: *Sonification, Semantic Web, Database, Query, SPARQL, luposdate, luposdate3000*

1 Introduction

For further analysis, purposes of debugging and sometimes teaching, data and its processing are often visualized. While high-quality visualizations often solve the investigated problems, there is still a demand for other ways of analysis by e.g. using another medium like audio [4, 19].

Sonification is typically defined *as the use of nonspeech audio to convey information [...] for the purposes of facilitating communication or interpretation* [15].

In a nutshell, the benefits of sonification approaches are typically the following ones: Sonification can help not only humans with visual impairments, but also help to reduce barriers for people with poor education. Audio can present multidimensional data in an easy-to-understand way, so that the most important aspects of the data and its processing can also be understood by people without science education. An additional medium also helps with learning new information [25].

We claim that these general benefits are of great value to the database community, too, increasing the attractiveness of and interest in database courses, technologies and scientific results in the area of databases, whenever there are sonification approaches for database technologies available. In this contribution, we propose a sonification approach of processing queries in a Semantic Web database. To the best of our knowledge and apart from our previous demo paper [9], this is the first approach to sonification of query processing in the literature.

Besides the above described general benefits of sonification, scientists and developers of database engines may detect anomalies by listening to the sonification of their database engine, which are hints for errors, and repeating sound patterns may be a hint for repeating (at least similar) calculations, which are to be optimized for increasing the performance. Hence sonification enables new ways of analyzing query processing for the purpose of debugging, performance tuning and research.

Our main contributions are:

- Sonification of query processing in a Semantic Web database,
- a flexible approach to map the data processed according to the query execution plan to sound effects,
- an easy-to-use web application[1] for visualizing the processing of Semantic Web queries along with its sonification, and
- an extensive user evaluation showing the benefits of sonification in the context of query processing in databases.

This contribution extends [9] by extensively introducing the basics of sonification in Section 2 and its related work in Section 3, and by the analysis of a user evaluation showing the benefits of sonification of database tasks in Section 8.

2 Basics of Sonification

There is a fundamental problem in sonification: There are currently no general standardized guidelines [22]. This allows a very creative development process, but makes the actual development of meaningful applications difficult.

As early as 1997 [16], however, the first considerations for the development of sonification applications were formulated, which are shown in Table 1.

Various experts published several recommendations based on these approaches in recent years, which are used as the basis for the conception of this work. The elaboration is primarily based on the *Sonification Handbook* [12], which describes in detail the basics of sonification.

In the Sonification Handbook, various interaction models are named that describe the possible implementations of sonification applications. With the *interactive data selection* the user is given a dynamic possibility to interact with the data, since (s)he can determine and change the subset of the data himself. There is also the *mapping interactions*, in which the user is able to set the mapping of data to the sound synthesis parameters in real time. Such an implementation is associated with some difficulties and often results in a higher degree of necessary prior knowledge on the part of the user. Operation must not be too complex, which is why it has to be simplified using simple control elements such as sliders, buttons or multiple option fields.

[1] Our demonstration with example queries, data and sonification mappings for all described use cases and target groups is available at `https://www.ifis.uni-luebeck.de/~groppe/soundofdatabases/`.

Table 1: The considerations of the first sonification report for successful sonification applications [16]

Concept	Description
Control	provides efficient, effective, accessible, parametric controls of the sound that constitutes the display medium.
Mapping	allows the design of new sonifications “on the fly” by giving the user flexible, intuitive control over which data dimensions control which sound parameter.
Integra-bility	facilitates data importation from a variety of formats to allow data from many different disciplines to be sonified.
Synchrony	allows easy integration with other display systems such as existing visual monitors, virtual reality systems, or assistive technology devices.
Experi-mentation	integrates a perceptual testing framework with the overall sound synthesis and mapping functions.

The Sonification Handbook also mentions three different models of sonification. A very simple model is the *non-interactive sonification* in which the user is not given the opportunity to interact with the application. Various experts are arguing against this approach, since at least a small form of interaction should be possible [12, 16]. In contrast, users influence in the *Parameter Mapping Sonification* approach the mapping of data to auditory dimensions, which is the primary interaction with the system. This model is used today in most applications of sonification [4, 12].

As for Parameter Mapping Sonification, *Model-Based Sonification* transforms the data into the auditory dimensions via a mapping, but the user cannot change this mapping her-/himself. The user interacts with the system via a virtual instrument. Model-Based Sonification is suitable for very large amounts of data with a high data dimensionality.

The Sonification Handbook also describes some guidelines that must be observed during development of sonification applications: Similar to how humans view new objects from several perspectives, a well implemented system should also have the option of several “sonification views”, i.e. several acoustic views. In this way, the same data can be viewed from different

Table 2: Different kind of mappings [12] from data features to sound synthesis parameters

Mapping	Description
one-to-one	Allows the strict mapping of one data feature to exactly one sound synthesis parameter.
one-to-many/ divergent	A data feature can be mapped onto various sound synthesis parameters.
many-to-one/ convergent	Several different data features can be mapped onto the same sound synthesis parameter.

perspectives. A statement that again emphasizes the importance of a flexible mapping of data to parameters. The different channels presented (visual and acoustic) should not be separated from each other, but used together and complement each other.

The mapping of data and parameters is formally described by a transfer function [12]:

$$Let\ \{\vec{x}_1, ..., \vec{x}_n\}\ with\ \vec{x}_i \in \mathbb{R}^d,\ and \tag{1}$$

$$g : \mathbb{R}^d \to \mathbb{R}^m,\ then \tag{2}$$

$$s(t) = f : \mathbb{R}^{m+1} \to \mathbb{R}^q \tag{3}$$

$$s(t) = f(\vec{p}; t) \tag{4}$$

$$s(t) = \sum_{i=1}^{N} f(g(\vec{x}_i), t) \tag{5}$$

If the d-dimensional input data from equation 1 are given with a parameter mapping function as in equation 2, a sound event is generated by a signal generation function (equation 3) which describes an acoustic q-channel signal as a function of time (equation 4). Here $\vec{p}$ is an m-dimensional vector of sound synthesis parameters of the signal generator. q is the number of different dimensions of the sonification. Finally a Parameter Mapping Sonification can be calculated (equation 5).

In practice there are primarily the *one-to-many*, *one-to-one* and *many-to-one* mappings, which are described in Table 2.

The Sonification Handbook provides four additional guidelines:

1. The sonification must reflect the object, its properties and relationships of the input data.
2. The transformation of the data must be systematic. A change in the data must be accompanied by a foreseeable change in the sonification.
3. The sonification must be reproducible with the same data.
4. The functionality of the sonification must be the same for different data sets.

A meta-analysis from 2013 [4] analyzed over 60 projects that dealt with sonification and examined which models were used. The analysis determined that a total of 30 different sound synthesis parameters were used for the mappings and that two methods were used particularly frequently for the sonification implementation: The tones are either generated by a simple synthesizer or a library of previously recorded tones is accessed. One conclusion of the meta-analysis is that sonification can only have a relevant benefit if the actual information and the relationships between the data are communicated in an understandable manner.

3 Related Work

Sonification deals with the transformation of data to sound. In doing so, the data, which are usually only shown graphically, are mapped in an additional dimension, thus opening up a further channel of perception for the viewer respectively listener. The sonification is intended to enable the listener to cope better with certain analysis tasks and to be able to concentrate his field of vision on certain events without other content no longer being perceived. One of the oldest examples of practical application is the Geiger counter, which enables the user to concentrate on the task at hand without taking his eyes off [25]. The research area of modern sonification is not new and exists for almost 30 years, if one considers the first ICAD conference in 1992 as the beginning of this movement. Since its publication, the research area of sonification has grown steadily. A search of all *Web of Science* databases in March 2019 showed that the use of the term *sonification* in literature has increased almost fourfold in the last 20 years [19]. However, a large-scale application of sonification in data analysis has not yet occurred and today's applications are increasingly limited to warning and information systems [19].

With the scientific data analysis not only the pattern recognition should be supported, but the scientific statements should be made accessible to a larger part of the population. This part is generally not familiar with the conventional analysis of data sets and also has no academic background [25]. In 2020 NASA published a series of videos that map some large objects or events in our galaxy [21] to music like a star cluster, in which the spatial arrangement of stars and their mass influence the generation of the sound.

The contribution in [1] deals with the ATLAS experiment, where ATLAS is a particle detector used in a particle accelerator. Until 2035 the ATLAS experiment will concentrate on the investigation of elementary particles such as leptons and quarks. It is obvious that the data obtained are highly complex and multivariable and that the experiments generate enormous amounts of data. [13] describes a transformation of this data to sounds, the aim of which is to investigate to what extent the notes could also provide inspiration for the musicians and whether the setting can also be of advantage in teaching.

Like the ATLAS experiment, many studies assume that, especially with large, multivariable data sets, the purely graphic data evaluation is tedious and that some correlations are not always immediately recognizable [5]. One work [25], for example, generates music from the multivariable data set of several thousand trees from five different conifer species, which are spread over 50 vegetation areas within Alaska. In total there are over 30 different variables per tree with the geographical latitude as the temporal dimension. The work concludes that it only takes a little time to explain the underlying assignment of the parameters to the auditory dimensions such as instrument, note and tone duration to the listener. Nevertheless, the sonification enables the rough distribution of the trees and the key message of the data to be made understandable without the listener having to look at and analyze the large amounts of data. Numerous other works and studies [24, 3, 11] show that the sonification in addition to the graphic representation helps the test persons to solve the tasks more accurately and with a better overall result. This gained accuracy goes hand in hand with a longer processing time of the tasks, since the test persons had to listen to the data sets several times in order to be able to make precise statements [24]. In their experiments, the test persons had no time limit for viewing the purely visual or listen to the auditory data.

In addition to the purely scientific application, sonification opens up teaching purposes. The additional sensory perception should help pupils at school and students to better understand the processes and relationships in a system and the algorithms used. It was shown that sonification is an alternative to conventional diagrams and graphs, which either improved the learning process or could at least be used as an equivalent alternative [27].

Artistic applications were also tested: For example, NASA, in collaboration with artists, published some videos that maps human efforts to explore the moon to music and thus compose a piece of music [26]. Various instruments are used for important events in space travel and the time sequences (months, years) for the rhythm.

New developments in the development of barrier-free applications through sonification can also be recognized. These applications enable visually impaired people to deal with visual data sets by mapping them to music [12]. There are already a large number of such applications that transform text into speech. However, applications that make more complex scientific data audible are a novelty. A prototype was developed that made 3D point clouds accessible and understandable for visually impaired people and examined which characteristic features could be recognized within these point clouds [3]. It turned out that the localization of objects and the determination of their size was possible for the test persons, but the geometric shape could not be recognized.

The paper [10] describes a sonification, where web tracking data is mapped in real time to sound in order to generate increasing awareness of the problem among users. The developed framework monitors the data traffic, filters connections to known web trackers, extracts particularly interesting events and sends them to an external audio framework using the Open Sound Control (OSC) protocol. The conclusion of the study is that sonification can make a significant contribution to sensitizing users to web traffic and the associated problems.

Sonifyd [14] is a multimedia and audiovisual environment to analyze images horizontally and vertically using a *scanner line* and to generate sounds based on a color-to-sound assignment. The aim of the work is increasingly directed towards the development of an application that, on the one hand, advances research in the field of sonification, but is intended to be used particularly in the context of art installations. For the development of Sonifyd the programming language *Processing*[2] and the *Siphon*[3] framework are used to transfer videos and pictures over Mac OS systems in real time. The MaxMSP framework is used to generate the sound using the OSC protocol.

A more practical application [18] focused on sonification on 3D audio in an interactive hypermedia environment. It is developed taking into account the design pattern “TheEarsLeadTheEyes” [2]. A simulation in Java has been developed, in which the users could navigate an avatar through a map. The Java Open Audio Library (JOAL) is used to make the tones playable in three-dimensional space. Based on a study the authors concluded that simple executions of sonification can be very effective. With increasing complexity, such as the occurrence of two different sound sources, the sonification led to increased confusion among the test subjects and thus to the opposite effect.

[2] https://processing.org/
[3] http://syphon.v002.info/

The software *Sonifigrapher* follows a conventional approach to sonification [23]. Light curves used to identify exoplanets have been converted into sound. It is possible to import these light curves as PNG files in order to sequentially convert the RGB channels of the image file into sound. The user has the option of fading out individual RGB channels and only adding audio to a certain area of the light curve. The audio programming language *CSound*[4] is used for the implementation. No study was carried out in this work, but the authors concluded that sonification can be used to highlight connections or correlations between variables.

At the moment there are no applications that allow access to the graphic representations of processes in a database.

4 Target Audience of Sound of Databases

With our sonification demonstration of a database, we address a large spectrum of target audiences and discuss the benefits of our sonification for these audiences in the following paragraphs.

General Public: Sonification is a means to offer a multimedia show to the general public at events attracting more people. By offering sonification applications to be used by any persons, sonification of databases may help scientists to increase the interest in database technologies in a playful way. It may help for an easy understanding of database technologies offering an audio-visual presentation of the data processing.

Students: Computer science students may more easily learn and don't forget the learned facts about database algorithms especially if these can be recognized in the sonification. For example, the sonification of merge joins may result in an ascending scale remembering the student that the input of merge joins must be sorted.

Our purpose is not only to reach students of computer science or related subjects but also to reach students of foreign subjects to give them an access to the basics of databases. So our main focus are associations like: What happens with my query request? How is it processed? How do I request special data? Sonification shall help them to recognize regularities and differences of varying optimizations of requests audio-visually.

To make computer science and its mostly theoretical topics more interesting to students and pupils, multimedia applications are beneficial. In cases that school students can experiment and create unique and interesting solutions on their own, it drums up interest for usually monotonous topics. Even in computer science there should be possibilities to arrange those topics in an exciting way. With the support of visual and auditive attractions pupils should playfully and perhaps musically learn the basics to work with databases. The main focus is learning with fun, to discover the features of such applications on their own to get access to databases and computer science topics.

Teaching Staff: Our demo is beneficial to teaching staff no matter if they are teaching at universities, schools or provide seminars for employees. Our solution provides an easy access to a complex topic to any level of students and pupils. Therefore different adaptations will be needed depending on the teaching method: presentation of results or experimenting. So in the first case the teacher will be able to record single query requests or create problem depending sound mappings in advance.

Developers of Database Engines: Irregular patterns and anomalies in the sound of the sonification are a hint for errors in the code, such that a sonification of a database engine helps developers to identify these errors. Furthermore, patterns in the sound may help to analyze performance issues and may provide hints which code should be optimized. Especially repeating sound patterns may be a hint for repeating calculations, which might be avoided improving the performance of the overall calculation. Scientists may also get ideas for efficient super operators replacing several simpler operators.

Visually Impaired: Experiments show that visually impaired can recognize the location of nodes in graphical data when using a proper mapping from the 3D location to auditory dimensions for sonification [3]. The graphical data are the operator trees of query execution plans for our proposed sonification. Hence our proposed sonification helps the visually impaired to understand the database technologies, especially query processing and its algorithms, and to make them audibly tangible for these technologies.

Musicians/Artists: All interested persons including musicians and artists are invited to experiment with a new tool to create unique sounds. We offer the option to use different instruments and sounds at the same time stimulating to assemble new rhythms and melodies.

5 The Semantic Web Database Engine LUPOSDATE3000

The Semantic Web database engine LUPOSDATE3000[5] [28] is the successor of LUPOSDATE[6] [6]. The focus of LUPOSDATE3000 is

[4] https://csound.com/

[5] Source code of LUPOSDATE3000: `https://github.com/luposdate3000/luposdate3000`

[6] Source code of LUPOSDATE: `https://github.com/luposdate/luposdate`

a modern architecture using latest database technologies (see Section 5.1) while going to support multiple platforms. We introduce the vision of a hybrid multi-model multi-platform (HM3P) database in [8], which spans over different platforms in operation and supports different data models in one single database, and its semantic variant - SHM3P database - in [7]. In future work we plan to develop LUPOSDATE3000 further to a SHM3P database. Currently, LUPOSDATE3000 is a semantic multi-platform database management system.

5.1 Internals and Sonification Web App

We develop LUPOSDATE3000 in the Kotlin[7] programming language in order to support different targets like the JVM, JavaScript and in the near future[8] native binaries for desktop, server, web and mobile environments. In this way it is possible to develop a common code basis for efficient database servers, distributed databases for cloud and IoT, mobile databases for operating on phones and tablets, and web applications for demonstrating and teaching purposes running completely in the browser.

LUPOSDATE3000 is optimized for modern multi-core cpus by introducing RDF-3X variants enabling partitioned input for efficient parallel [28] and distributed processing. Our next steps in future work include the full-fledged support of Internet-of-Things (IoT) environments.

LUPOSDATE3000 uses a dictionary to map all values to integer ids. Dictionaries are decreasing the memory footprint during query processing as well as reducing the space used for indices. For the sonification, the integer ids can be directly used for manipulating auditory dimensions like the pitch or duration of the sound.

These ids are then stored in an index similar to RDF-3X [20] using all 6 collation orders for maximizing the use of the fast merge joins on pre-sorted data retrieved from the index. The sonification of processing pre-sorted data generates regular sound patterns with a distinctive sound experience. Operators, which do not generate their output in a sorted way any more (like duplicate elimination using hash tables), also do not generate an ascending scale, which is immediately recognized by users of the sonification. Hence each operator generates a unique sound pattern during query processing resulting in a complex sound ensemble for complex queries.

During query evaluation, LUPOSDATE3000 uses both column and row iterators, preferring the column iterators where applicable. The sonification of a single column value generates a single tone, whereas the generation of a complete row can include the sonification of multiple column values. Hence the difference between the usage of column and row iterators would be audible, too, but the sonification in LUPOSDATE3000 currently only supports the sonification of a single column by implementing column iterators.

The *sonification software*[1] offers a single page application in the web, which provides access to a LUPOSDATE3000 database server, but also fully integrates the LUPOSDATE3000 web version, such that the sonification web application can run stand-alone in the browser. We use Tone.js[9] as simple-to-use library for generating tones with the add-on tonejs-instruments.js[10] for using instruments, and vis.js[11] and vis-network.js[12] for visualizing operator trees and animating query processing.

6 Sonification of Query Processing

Our proposed sonification of query processing generates a tone for each exchange of intermediate solutions[13] between the operators in the operator tree of the query execution plan (see Figure 1). The sequence of intermediate solutions determines the sequence of generated tones. For the purpose of collecting when which intermediate solutions are exchanged between which operators, we introduce a new kind of operator, which we call *log operator*. We insert log operators between each pair of operators connected by an edge in the operator tree (see Figure 2). In this way we log each step of processing queries in LUPOSDATE3000 for sonification and also for debugging purposes. Please note that during normal operation, the log operators are left out for maximal performance.

Most sonification approaches use the following auditory dimensions (out of 30) [4]: pitch, loudness, spatialization, duration, brightness, timbre, tempo and spectral power. These auditory dimensions are very useful for a mapping of query processing, too. Hence we need quantitative numbers for the mapping of query processing to auditory dimensions during sonification. We first normalize these quantitative numbers to the $[0;1]$ domain and afterwards scale it to rational minimum and maximum values of the auditory dimension to be mapped to.

During query processing, the most relevant

[7] `https://kotlinlang.org/`

[8] After the introduction of a new memory manager including a new garbage collector for Kotlin/Native, such that current restrictions are lifted (see `https://youtrack.jetbrains.com/issue/KT-42296`).

[9] `https://tonejs.github.io/`

[10] `https://github.com/nbrosowsky/tonejs-instruments`

[11] `https://visjs.org/`

[12] `https://github.com/visjs/vis-network`

[13] To be more precise, for each of its column values.

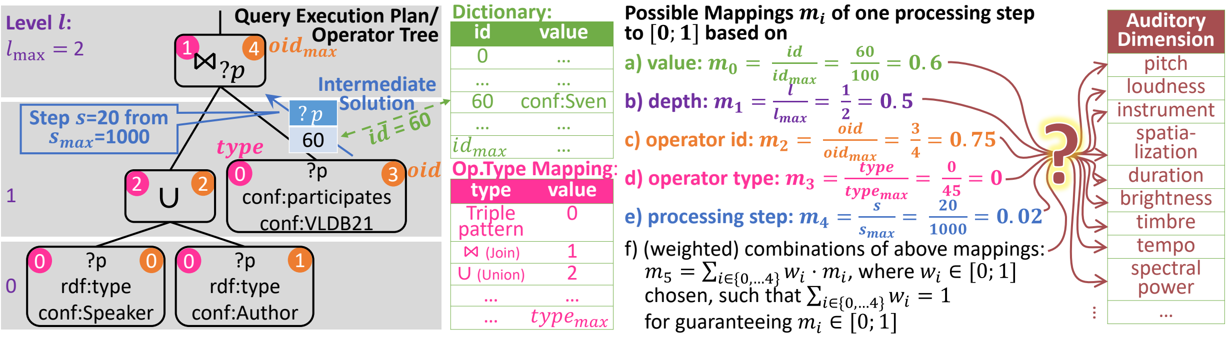

Figure 1: Query execution plan of a simple query with an intermediate solution between two operators along with possible mappings of this processing step

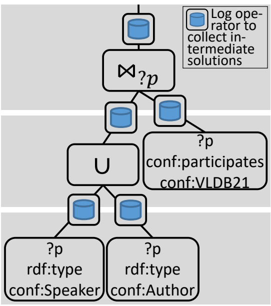

Figure 2: Operator tree to collect intermediate solutions of each query processing step

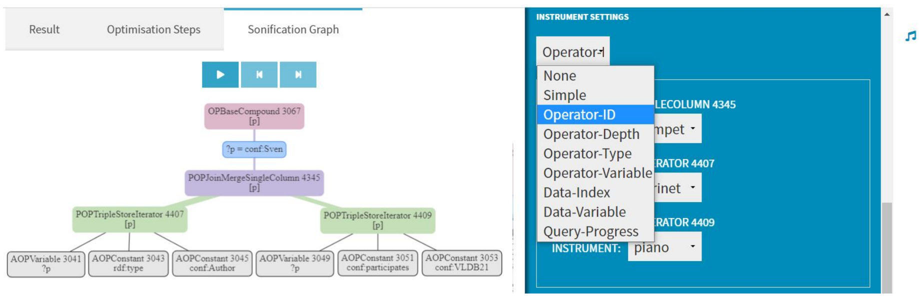

Figure 3: Sonification web app with flexible configuration possibilities: The auditory dimensions pitch, instrument, loudness, spatialization (left to right speaker), duration, melody (e.g., dur cadence and hit formula), chord and octave can be configured in a flexible way to be mapped from operator id, type, depth (in operator tree), variable ("Data-Variable"), variable combination ("Operator-Variable"), bound value ("Data-Index") and query progress. The merge join operator just determined the intermediate solution "?p = conf:Sven".

information are the sequence of intermediate solutions consisting of bindings of values to variables, which are the output of certain relational operators serving as input of the succeeding operator in the query execution plan (see Figure 1). Because LUPOSDATE3000 uses a dictionary, the values of the intermediate solution are already represented as integer numbers for which a maximum value can be determined. In this way we have a scale for the mapping to the quantitative auditory dimensions.

Overall, when considering query execution plans and query processing, the following information provides quantitative numbers during query processing (see Figure 1):

- The integer id representations of the values in intermediate solutions,
- the ids of the operators, their depths in the operator tree of the query execution plan or their types having the intermediate solution as input or output, and
- the temporal aspect, i.e., at what time relative to the whole query processing the intermediate solution has been generated.

While the mapping to the auditory dimensions can be configured in a flexible way, according to our experiences the most useful ones seem to be:

- for understanding the operator algorithm (or at least for recognizing characteristics of its in- and output): a mapping from the intermediate solution values to the pitch and the operators - between which the intermediate solution is exchanged - determines the used instrument, such that operators can be easily distinguished from the other (see Figure 3), and
- for generating inspiring melodies and sounds (not only for musicians/artists): a mapping from the operator depth to pitch and different operators may use different instruments.

Dependent on the query and the data, other mappings may generate convincing sound effects, too, such that users are encouraged to extensively experiment with the mapping configuration.

7 Demonstration

We describe the different used queries for our sonification demonstration. The used queries are presented with an increasing complexity and we motivate their use in the demonstration by discussing their purposes like learning effects of students.

7.1 Sonification for Guessing Single Operators

Experienced database developers and scientists may be asked for their guess what kind of relational operator is introduced in the sonification demonstration for the following queries. Novices to database technologies may be asked about their guess of the functionality of the considered query.

One Triple Pattern: The first query for the sonification demonstration is a simple query consisting only of one triple pattern. The listeners of the sonification hear an ascending scale getting to know that the solutions of triple patterns can be determined by accessing an index of pre-sorted data.

Filter: The second query adds a filter. Listeners of the sonification recognize that not all intermediate results are passing the filter operator playing the previous tone generated by the index scan with another instrument.

Merge Join: This query consists of a binary operator having two inputs and one output. Two index scan operators as well as the output of the merge join play with different instruments. Listeners notice that only if the two inputs of the index scan operators are the same, an output is generated.

Optional-Clause/Left Outer Join: This query contains an optional-clause, such that a left outer join occurs in the operator tree of the query execution plan. In comparison to the merge join, the input of the left index scan passes always the left outer join - even if it is not combined with the input of the right index scan.

Hash Join: This query consists of a merge join of two triple patterns and a succeeding hash join combining the result of the previous merge join with a triple pattern. The hash join is implemented as pipeline breaker by generating the output after reading the complete input of one operator.

7.2 Sonification of Complex Queries

In this section we describe the sonification of more complex queries allowing further analysis and generating a more complex ensemble of sound effects.

Recognizing Single Operators in Complex Queries: Listeners of sonifications of very complex queries may recognize single operators in the complex ensemble of different operators and their generated sounds during query processing.

Common Subexpressions: Queries consisting of the same subexpressions several times generate a repeating sound during sonification, such that listeners may notice these common subexpressions.

Similar Queries: The sonification of similar queries generates similar sound with some few differences (like passing other solutions in different filters), which might be recognized by the listeners.

Queries generating inspiring Sounds: Some sonifications of complex queries may generate interesting regular sound patterns, which may inspire musicians and artists for new musical compositions.

8 Evaluation

In total, we propose and test five hypotheses for sonification of databases:

Hypothesis A: Sonification has a positive effect on the analysis of database data and pattern recognition. In order to confirm this hypothesis, the effect of sonification on general analysis tasks must be examined. In the special context of databases, for example, it is examined whether tasks by sorting data within a database can be better solved with sonification. It should also be investigated whether certain properties of the used allocation of information and sound synthesis parameters can be learned indirectly without the allocation being specifically explained. This should also enable the subjects to discover anomalies in the normal behavior of the database.

Hypothesis B: Sonification helps to make scientific data more understandable for non-academic people. In order to answer the hypothesis, it must be specifically considered to what extent non-academic subjects perform compared to subjects with an academic background. It is not to be expected that test subjects without an academic background will ultimately show better results. Rather, it is to be expected that a significant improvement in the results will occur if the visual data are also mapped to music. This improvement should be greater than that of the test persons with a scientific background, since it can be assumed that this test group can interpret the purely visual data much more easily. In addition, a tendency should be observed that the test subjects develop an increased interest in databases with sonification.

Hypothesis C: Sonification is also suitable in the context of databases for visually impaired people. It is not to be expected that a sufficient number of test subjects with visual impairments will take part in the planned online survey in order to be able to make a statement about hypothesis C. Instead, after several tasks in which visual and acoustic information are presented, the test subjects should work on a task with purely acoustic information. It is to be expected that the previous tasks will have a learning effect and that the test subjects will be able to work successfully on the tasks despite the lack of a visual channel. It is assumed that visually impaired people could then also solve these tasks.

Hypothesis D: Sonification has a positive learning effect for pupils at school and students. Due to the university background of this work, the investigation of this hypothesis will primarily be related to students. It is to be investigated whether sonification can contribute to the fact that unknown facts can be learned more easily through sonification. If the learning effect is positive, it can also be assumed that the sonification generally leads to a better understanding and an increased interest in databases and the sonification itself.

Hypothesis E: Sonification can be used to create artistic works from data. In order to confirm this hypothesis, it should be examined how exciting and interesting the sonification is. If interest in general can be aroused, it can be assumed that corresponding artistic multimedia applications can also arouse the interest of a larger group. The test persons are specifically asked for their assessment of whether the sonification favors artistic applications.

In order to examine the hypotheses raised, an online survey was processed.

8.1 Online Survey

We select a suitable platform for the online survey, which provides the necessary tools for the investigation of the hypotheses. The test subjects' user interface must be kept simple and clear. This ensures that the test subjects are not distracted by other content in the browser during the survey. The platform must offer the possibility of freely choosing the arrangement and number of the various elements. In the online survey, the subjects have to watch various videos and answer questions, such that a video integration is a prerequisite for the questions. The platform must support layouts from empirical user evaluations [17] and must be freely available. The results of the survey must also be freely available and access to them must not be restricted. The individual answers of the test persons must be viewed and analyzed individually. The mentioned criteria are met by the Google forms[14] platform.

[14] https://www.google.de/intl/de/forms/about/

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX ub: <http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.
      owl#>

SELECT ?student
WHERE {
 ?student rdf:type ub:Student .
}
```

(a) SPARQL Query

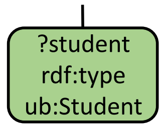

(b) Operator Tree

Figure 4: SPARQL query and operator tree of the first task of the online survey

8.1.1 Tasks in the Online Survey

The survey consists of a total of four tasks, with two videos being shown for almost every task. Each of exercises 1 to 3 first shows a video of an operator tree without audio and then the same video with audio. After each video, the subjects have to answer a series of tasks. In the fourth task, the test subjects are shown a video with only acoustic information. Here it is examined which unconscious learning effects the test subjects achieve in the previous videos and can apply them at this point.

In the following, each individual tasks will be explained in more detail.

1. Task: Simple operator tree
For the purpose of a simple introduction, the operator tree chosen for the first question is very simple and only includes a triple store access.

After both videos, the test subjects are asked a question about the sorting of the data: "What did you notice about the sorting of the data?". The data is not sorted according to the lexicographic order, but rather according to their index value in the database dictionary. Accordingly, no sorting is recognizable for the test person in the video without audio. Sorting based on the ascending scale can only be recognized in the case of the second video with audio.

The SPARQL query is shown in Figure 4(a) and the associated operator tree is shown in Figure 4(b).

2. Task: Operator tree with increasing complexity
For the second task, a more complex operator tree is used, which consists of three triple store accesses and two join operators. The text within the join operators

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX ub: <http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.
     owl#>

SELECT ?student ?course
WHERE {
 ?student rdf:type ub:Student .
 ?student ub:takesCourse ?course .
 ?course ub:name "Algorithms" .
}
```

(a) SPARQL Query

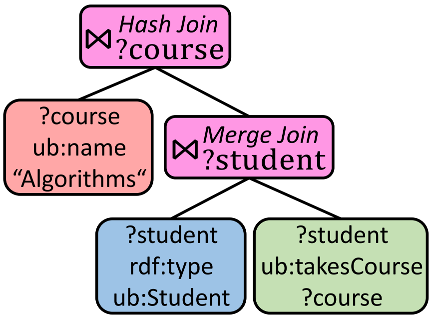

(b) Operator Tree

Figure 5: SPARQL query and operator tree of the second task of the online survey

was made unrecognizable for the test persons. Again two videos are shown, whereupon a few questions are asked. In addition to the familiar question from the first task (sorting), the question “The tree consists of several edges. Which was the most active?” is asked. The test persons should indicate over which colored edge most of the data flow. A task that is easy to solve with the purely visual representation, as the edge thickness also increases dynamically according to the activity. The X position of the operators was transferred to the stereo loudspeaker channels. In addition, a mapping with regard to the duration of the tone is selected. Data values with the variable `?student` have a significantly shorter tone than data values with the variable `?course`. Join operators use the guitar as an instrument, while all other operators use the piano. These mappings are not communicated to the test persons and should be learned unconsciously. Figure 5(a) shows the SPARQL query and Figure 5(b) the associated operator tree.

3. Task: Operator tree with high complexity
In the third task, the operator tree is more complex with an additional join and triple store operator. All previous mappings are retained. In the third task, the same questions are asked about the videos as in task 2.

The figure 6(a) contains the SPARQL query and Figure 6(b) its operator tree.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX ub: <http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.
     owl#>
PREFIX instance: <http://www.Department0.University0.edu/>

SELECT ?student ?course
WHERE {
 ?student rdf:type ub:Student .
 ?student ub:takesCourse ?course .
 ?student ub:advisor instance:HiWi1 .
 ?course ub:type ub:course .
}
```

(a) SPARQL Query

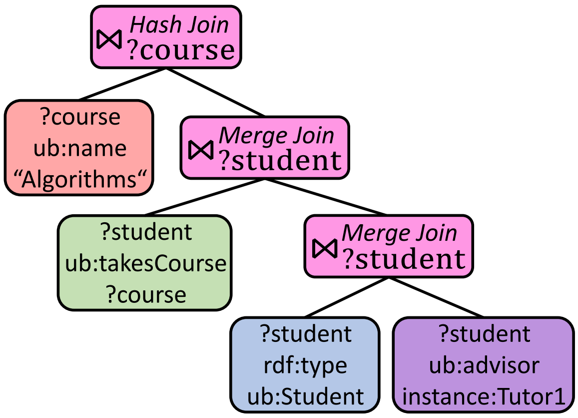

(b) Operator Tree

Figure 6: SPARQL query and operator tree of the third task of the online survey

4. Task: Only acoustic information
The last task only contains an audio file that plays the acoustic information of a simple operator tree. The known mappings are retained. A new sort operator, which sorts the data in descending order according to the `?student` data, is placed at the end of the operator tree. This operator uses a new instrument: the xylophone. You can hear through the strictly descending scale that this operator puts the data in a different order. This task tests whether the previous mappings were learned unconsciously and whether they can be used in this task. The familiar questions about the sorting and activity of the edges are asked again. It is checked whether the test subjects can correctly recognize the three operators. The middle operator appears to be the most active, since most of the tones “in the middle” are emitted from the stereo speakers. In addition, the question “Was there a join operator?” checks whether the mapping of the guitar to the join operators is recognized and also unconsciously perceived as such a mapping. Finally, the subjects are

```
PREFIX ub: <http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.
      owl#>

SELECT ?student ?course
WHERE {
 ?student ub:takesCourse ?course .
 ?student ub:name "Algorithms" .
} ORDER BY DESC(?student)
```

(a) SPARQL Query

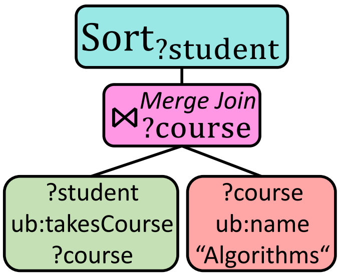

(b) Operator Tree

Figure 7: SPARQL query and operator tree of the fourth task of the online survey. The operator tree is hidden to the test subjects.

asked whether "[...] there is a new operator that has not been shown before?". The new instrument should make its existence audible. If the answer is positive, the test subjects are asked about the function of the new operator.

After the questions in the fourth task, the test subjects are briefly explained which function the join operators have. It will later be checked whether the sonification will help retain this new information.

Figure 7(a) shows the SPARQL query and Figure 7(b) the resulting operator tree for the fourth task.

8.1.2 Additional Questions in the Online Survey

At the end of the first part of the online survey, the test subjects are asked to rate seven different statements using a 5-point linear scale. The 1 represents the evaluation "strongly disagree" and the 5 "definitely agree".

The seven statements to be rated are:

A) The sonification of the query processing provides me a better overview
B) The sonification of the query processing helped me to better answer the questions asked
C) The videos with the sonification are more interesting than the ones without the sonification
D) Due to this survey I am more interested in databases
E) I find the idea of sonification of query processing exciting
F) Due to the presented sonification I am able to remember better what the function of the join operator is
G) I can imagine that artists could use this sonification to 'compose' new music

The online survey[15] was sent to the student mailing list of the University of Lübeck, and was distributed via social networks, messenger services and via word of mouth. The aim was to reach as many test subjects as possible, both with and without an academic background.

8.2 Results of the Online Survey

The online survey ran over a period of four weeks (May 13th, 2021 - June 9th, 2021). During this time, 31 different test subjects took part in the online survey. In addition to the questions described above, the subjects were asked about their age, highest educational qualification and knowledge in the field of computer science. The acquired data was recorded anonymously. IP addresses were not saved.

To answer the hypotheses from section 8, a total of four test groups were defined:

Group A: Subjects without academic background: 17 subjects
Group B: Subjects with academic background: 14 subjects
Group C: Subjects who have no knowledge of computer science: 9 subjects
Group D: Subjects who have completed training/studies in computer science or have otherwise acquired basic knowledge: 22 subjects

The individual test groups are not disjoint.

8.2.1 Results of the 1st Task

The results of the first task are shown in Figure 8.

In the first task, the data was not sorted lexicographically. It was therefore expected that the majority of the test persons would choose the answer *the data were not sorted at all*.

Overall, it can be seen that in addition to the expected answer, a large proportion selected the answer *The data are sorted according to a criterion that I cannot see*. Accordingly, 41.9% initially selected the answer *sorted by unknown criterion*, 51.6% *not sorted* and 6.5% *I don't know*. Hence many test persons expected some sort of sorting based on the question, but could not identify any sorting criterion. Therefore the answers *sorted by*

[15] https://forms.gle/5GyNr77X3Wsn4byHA

unknown criterion and *not sorted* are rated as correct in retrospect.

A strictly ascending scale can be heard in the sonification video, as the data has already been pre-sorted according to its index. Therefore it was assumed that the majority of the test persons would choose the answer *strictly sorted* for this task. In fact, only 26.4% respondents chose the correct answer. 6.5% chose *mostly sorted, even if some elements are out of line*. Despite the ascending scale, 44.5% of the test persons stated that the data was sorted according to an *unknown criterion*. In the retrospective, it can be assumed that the test subjects continued to rely on visual impressions, i.e. the values within the data nodes, when choosing a criterion. Only 16.1% indicate that the data is *not sorted*.

Looking at the results of the different test groups, it becomes clear that both groups A and C have a significantly stronger learning effect, as the number of correct answers in group A increases by 35.29 percent and in group C by as much as 55.56 percent. Groups A and C thus improve significantly more than groups B and D. However, with an increase of 21.43 and 27.28 percent, these two groups indicate that there is a sorting, they just do not perceive them as strict.

It is striking that, despite the sonification, groups B and D increasingly chose the answer *not sorted* (B: 28.57%, D: 22.73%), while this answer with the sonification of groups A and C was not or only poorly selected (A: 5.88%, C: 0%). This confirms that groups B and D, due to their academic background and familiarity with subject-specific tasks such as data sorting, tended to rely on visual impressions and not purely acoustic information.

In groups A and C, there was also no person in the video without sound recording who chose the answer *I don't know*. At the beginning, groups B and D had a greater degree of uncertainty to make concrete statements (B: 14.29%, D: 9.09%) than groups A and C.

8.2.2 Results of the 2nd Task

Sorting the data: Figure 9 presents the results of the online survey of task 2a. In the second task, the data is presorted for the join operation. The data is selected one after the other by the triple store accesses based on the ID. However, if the data values are only observed according to their temporal appearance, then these appear arbitrary. With knowledge of the functionality of join operators and sorting algorithms, it can be assumed that groups B and D will assume a *strict sorting* or will regard the data as *mostly sorted*. For groups A and C, a selection of *sorted by unknown criterion* or *not sorted* should be chosen accordingly.

This preliminary consideration has mostly turned out to be correct, since groups B and D have for the most part chosen the answer *strictly sorted* (B: 57.14%, D: 45.45%), while groups A and C mainly choose the answers *sorted by unknown criterion* or *not sorted* (both options in total for A: 47.06%, C: 44.44%).

Due to the pre-sorting of the data for the join operators, an ascending scale can be heard during playing the sonification, even if the accessed data does not always have a strict index sequence. The new instrument for the join operators (guitar) also has a deeper sound than the piano. The increasingly higher tone should help all test subjects to recognize the general sorting. This was also the case for groups A and C, even if many test persons again chose the answer *sorted by unknown criterion*. Here it can again be assumed that a type of sorting has been recognized, but the lack of visual representation led to this choice. If you put these two answers together, group A recognizes with an increase of 23.53 percent and group C with an increase of 33.34 percent that there is a sorting, but the data is not strictly sorted. For groups B and D the number of answers *strictly sorted* decreases by 21.43 and 22.72 percent, while the answers for *sorted according to unknown criteria* and *mostly sorted* together by 14.29 and 22.72 percent increase.

It is interesting to observe that in the three groups A, B and D the uncertainty increases with the sonification or at least remains the same (A: 11.76% $\rightarrow$ 17.65%, B: 14.29%, D: 9.09% $\rightarrow$ 13.64%), while the uncertainty in group C decreases slightly (22.22% $\rightarrow$ 11.11%).

Overall, the sonification helps to identify a sorting in the event of a lack of expert knowledge. These tendencies can also be seen in groups B and D, even if the sorting recognized in advance was usually only specified more precisely after hearing the sonification.

Recognizing different operators: The results for task 2b are shown in Figure 10. As shown in Figure 5b), the operator tree has three triple store accesses, whereby the left operator was colored red, the middle one blue and the right one green. The test subjects are asked to determine over which colored edge of the triple store access operator most of the data pass. For this task, the visual information was already sufficient to solve the task. The solution to the task is additionally simplified by the dynamic edge thickness. Therefore it was assumed that all groups choose the answer *green*. This assumption has largely been confirmed. A total of 87.1% of the test persons chose the answer *green*. It is unexpected that 9.7% decided for the answer *blue* and 3.2% for the answer *All equally active*. Due to the low number of these wrong answers (4 out of 31), it is assumed that these test persons did not understand the task correctly, that technical problems such as insufficient resolution or

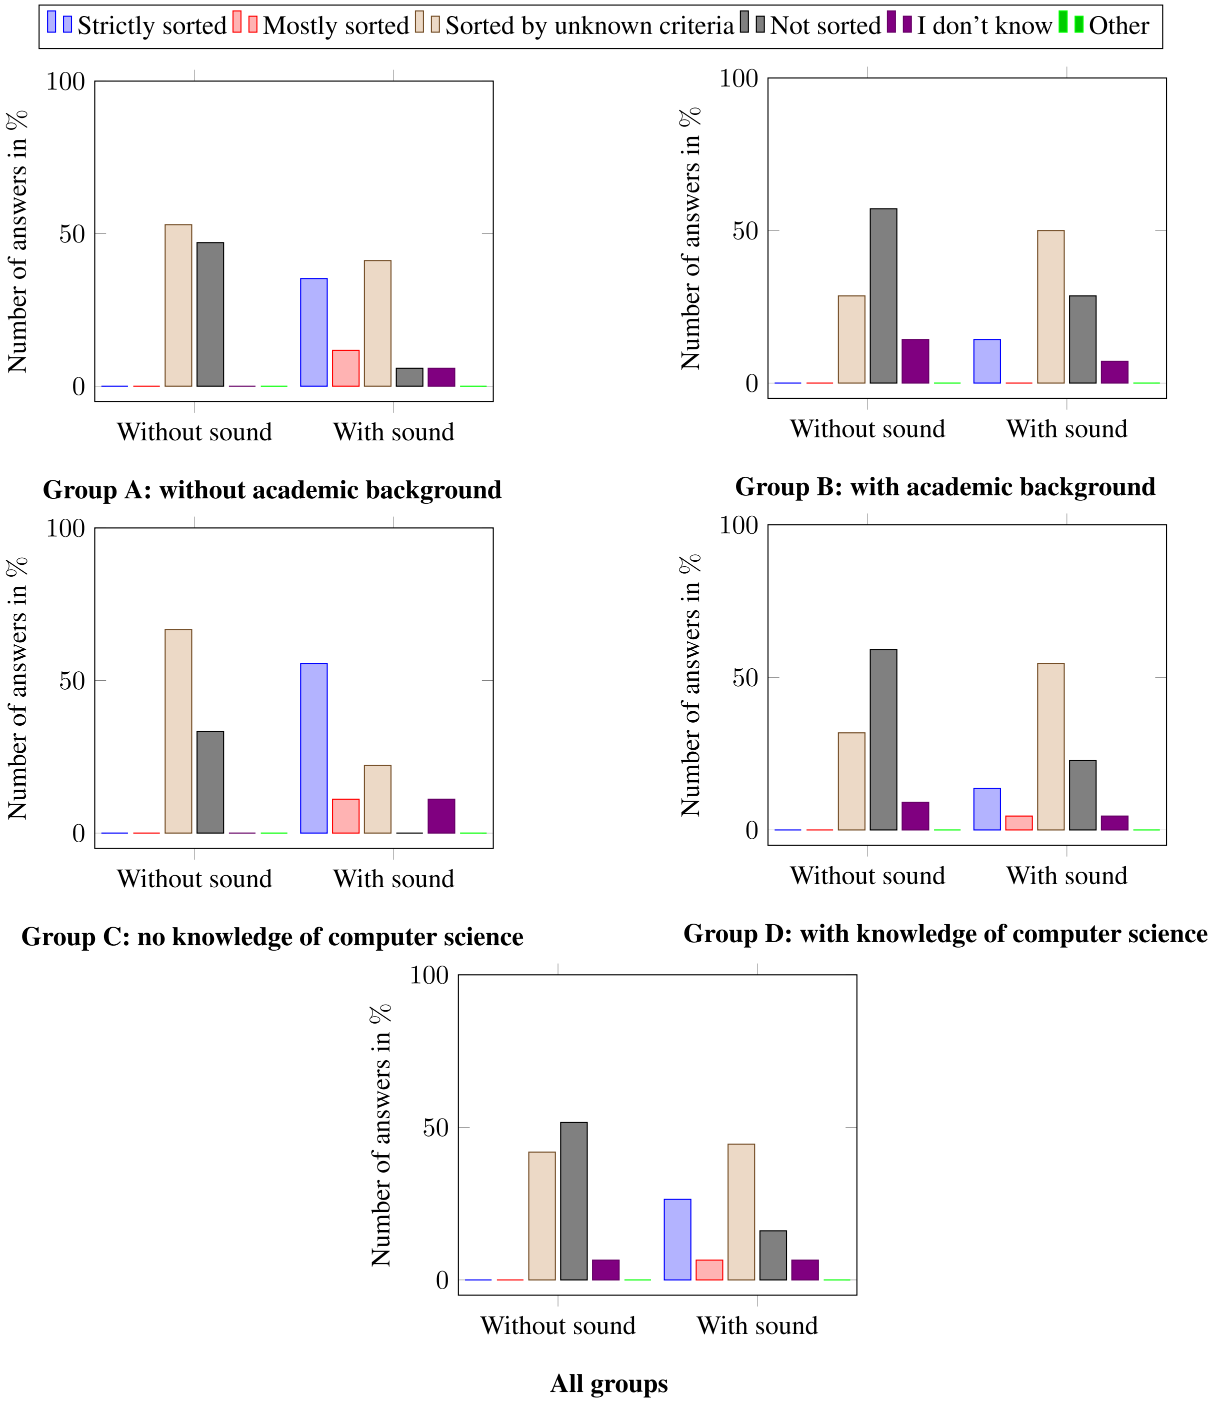

Figure 8: The results of the first task (sorting the data) for all groups. The results are shown without sound on the left and with sound on the right.

lack of concentration due to the increased complexity of the operator tree of the correct solution to this task occurred have distracted.

However, a slight improvement can be seen after playing the sonification. The horizontal position (x-axis) of the triple store access operators was mapped to the left and right speaker channels. This mapping was not explicitly communicated. Overall, the number of correct solutions increased to 96.8% after playing the sonification, which is an improvement of 9.7 percent.

This increase can be seen in all test groups, while group C had chosen the correct answer for both videos

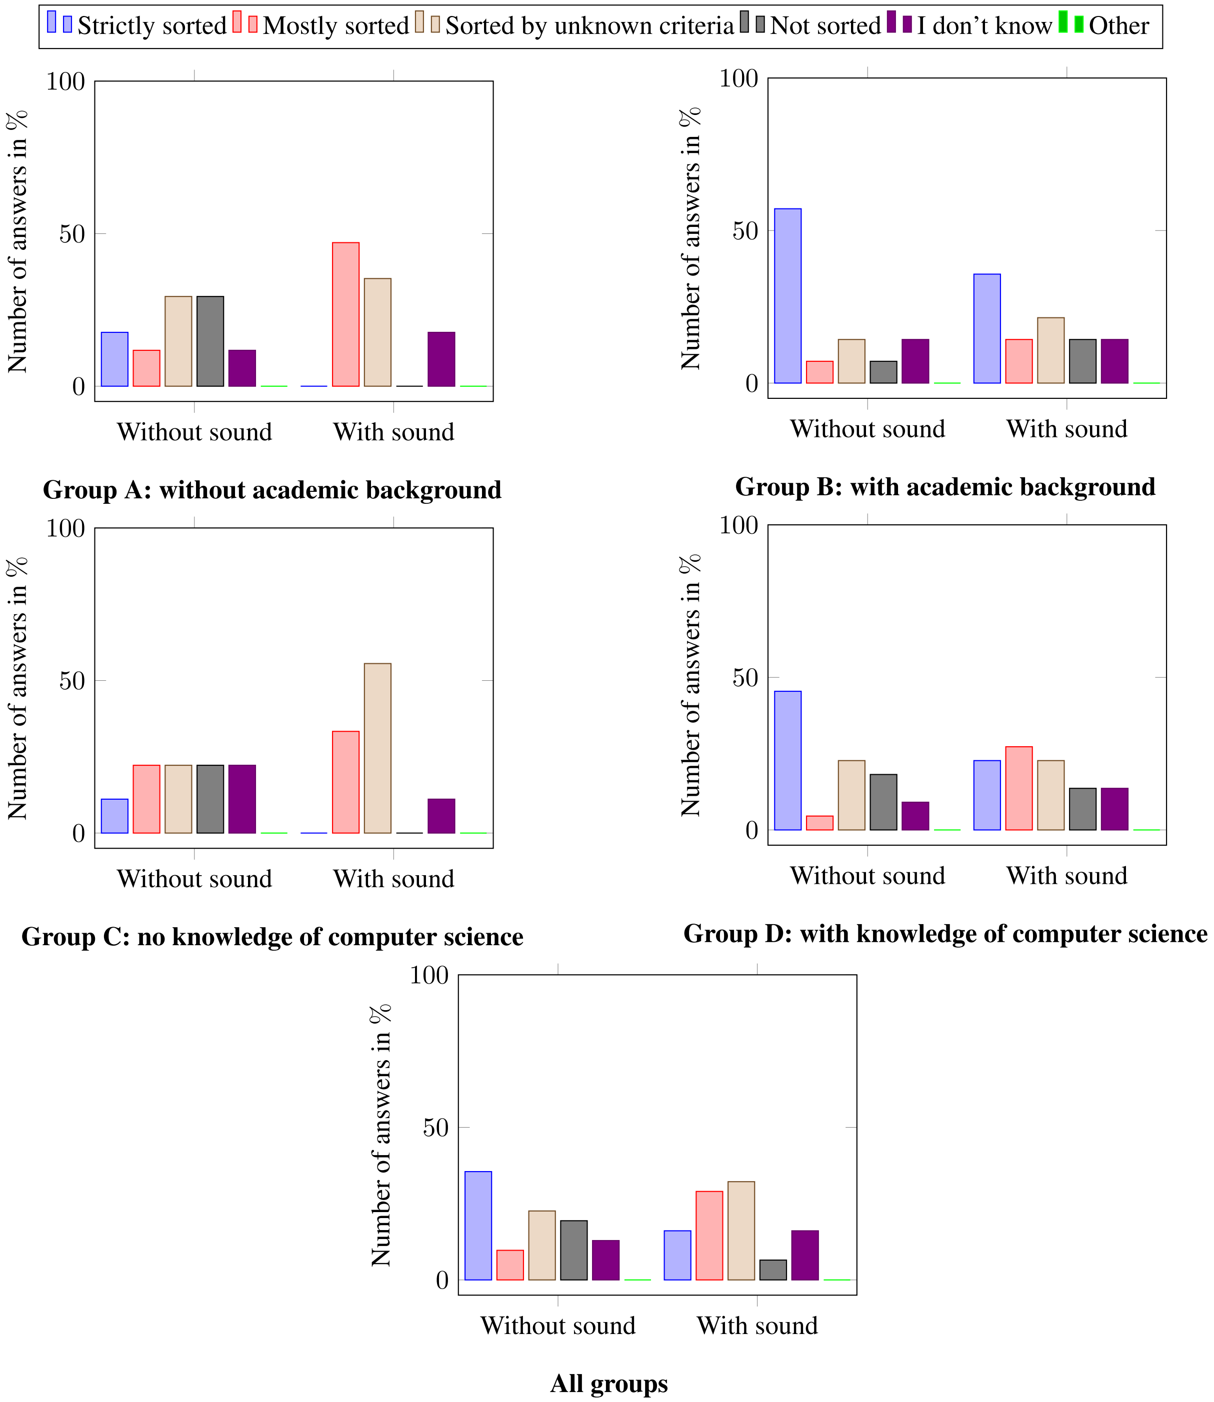

Figure 9: The results of the second task (sorting the data) for all groups. The results are shown without sound on the left and with sound on the right.

at 100%.

8.2.3 Results of the 3rd Task

Sorting the data: In the third task, an additional triple store access and join operator was added, which increases the complexity of the operator tree. Overall, fewer data ran across the operator tree in this animation, which is the reason for sorting is much more difficult to identify. Sorting is very difficult to recognize from a pure visualization. For this task it is to be assumed that a large part of the test persons now choose the *Unsorted* answer or the increasing complexity can be seen in an

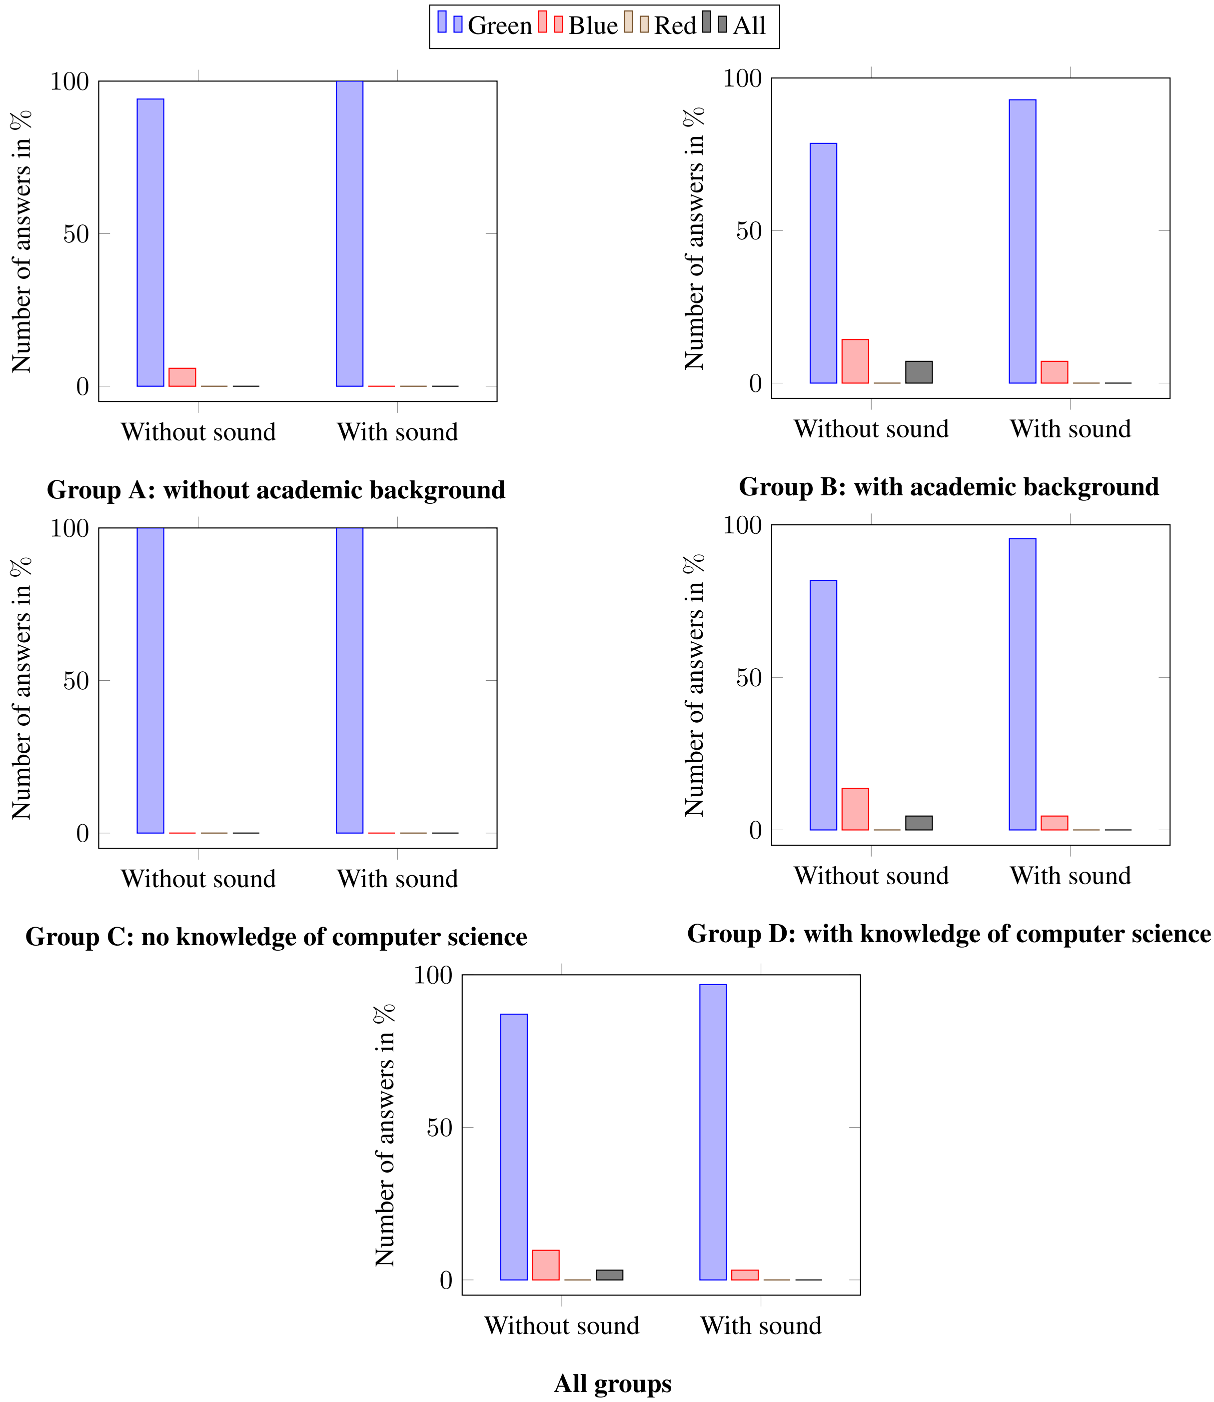

Figure 10: The results of the second task (recognizing different operators) for all groups. The results are shown without sound on the left and with sound on the right.

increasing uncertainty due to the choice of the *I don't know* answer. The correct answer is *sorted according to unknown criteria* or *mostly sorted*.

The results for task 3a are presented in Figure 11. All in all, this preliminary consideration is confirmed, since the number of answers from *not sorted* outweighs with 35.5%, followed by the answer *sorted according to an unknown criterion* with 25.8%. In the third task, the subjects probably assumed an expectation that the data would be sorted in some way.

It is unusual that a relatively large amount of 22.22% from group C has chosen the option *strictly sorted*. It is possible that the increasing complexity has led to excessive demands and therefore to this choice.

Similar to task 2a, a tendency from low to high tones are recognized for the sonification, whereby individual data tears out of this scale significantly more often. Accordingly, it was to be assumed that the answers *mostly sorted* and *sorted according to an unknown criterion* would be chosen. The assumption is confirmed: The answers that the data is *not sorted* decreases sharply with $35.5\% \rightarrow 16.1\%$, whereby these changes occur mostly for groups A and C (A: $41.18\% \rightarrow 5.88\%$, C: $22.22\% \rightarrow 0\%$). This results in an increase in the response *strict sorting*, *mostly sorted* and *sorted according to an unknown criterion* for both groups, with the latter increasing the most.

In groups B and D, a similar reduction in the *not sorted* answers can be seen along with an increase in the three “sorted” answers, even if the change is much more moderate than in groups A and C. This also correlates with the uncertainty of the groups. While the *I don’t know* answers from groups A and C drop to 0% for sonification, they remain the same for groups B and D.

It is also interesting to mention that in groups A, C and D the answers to *sorted according to an unknown criterion* increase by an average of 18.36 percent, while in group B they remain the same.

Overall, the sonification seems to help to identify a sorting where it can no longer be seen purely visually. However, a greater degree of uncertainty remains in this task.

Recognizing different operators: We present the respondents’ answers to task 3b in Figure 12. For this task, the selected mapping from task 2 is still used. Due to the additional triple store access and the lower number of data to be processed, the determination of the most active operator is more difficult, but should be trivial due to the dynamic edge thickness. In addition to the operator tree from exercise 2b, a purple operator is added to the right in the tree. It is to be expected that the majority of the answers would fall to *green*.

Even if, as expected, the majority of the test persons stated *green*, some of the answers were *blue* 9.7%, *purple* 6.5% and *all* 12.9%. Group B chose the answer *All* the most with 21.43%.

Although the differences between the different operators are small, especially in the third task, this deviates significantly from the assumption.

The sonification no longer leads to an improvement in the third task. Overall, the number of correct answers *green* even decreases slightly by 6.5%, while the answers to *blue* increase by 9.7 percent.

Only the renewed increase in complexity comes into question for the general deterioration. Given the many sound synthesis mappings, the sonification seems to lead to confusion for some test subjects and not to help. Nevertheless, it has to be noted that the majority of the test persons had already chosen the correct answer with the purely visual video.

8.2.4 Results of the 4th Task

Sorting the data: We present the results of task 4a in Figure 13). In task 4a we play only one audio file without any visualizations. Here two triplestore access operators use the piano instrument, the operators of which are connected to a join operator playing the guitar instrument. A sort operator, unknown to the test subjects, is added above the join operator. This sorts the data in descending order using the xylophone. The data of the triplestore access operator and the join operator operate with strictly sorted data, which could be heard on the ascending scale. It is assumed that the previous videos caused an unconscious learning effect with regard to the connection between the sorting and the pitch, which is the reason for a probably recognition of the present sorting by the test subjects. It is expected that the test subjects would mainly choose the answer *strict sorting*, while some could change their choice to *mostly sorted* by sorting in descending order using the sort operator.

It can be seen positively that overall neither the answer *not sorted* nor *sorted according to an unknown criterion* was chosen. However, the majority of the test persons (64.5%) opted for *mostly sorted*, while only 19.4% stated that the data was *strictly sorted*. Group C has the largest share of 28.57% of *strict sorting* answers compared to the other groups.

The sudden descending sorting by the sort operator will probably have prevented the majority of the test persons from choosing the *strict sorting* answer.

The exclusive acoustic channel unsettled a small number (5, 16.1%) of the subjects. The uncertainty was lowest in group B with 11.11%.

Overall, the hoped-for unconscious learning effect in this task can be confirmed, after all, the majority of all test subjects were able to recognize a sorting, even if no visual information was available at all.

Recognizing different operators: In the fourth task, the x-position of the operators is also mapped onto the stereo loudspeaker channels. If you listen carefully, you can hear that there are a total of three different positions: far left, center and far right. The triple store access operators that use the piano can only be perceived on the left or right, while the join and sort operators are in the middle.

The test persons should recognize the number of these three operators and indicate for which of these operators a particularly large number of tones are played. The localization used is not mentioned and could only be

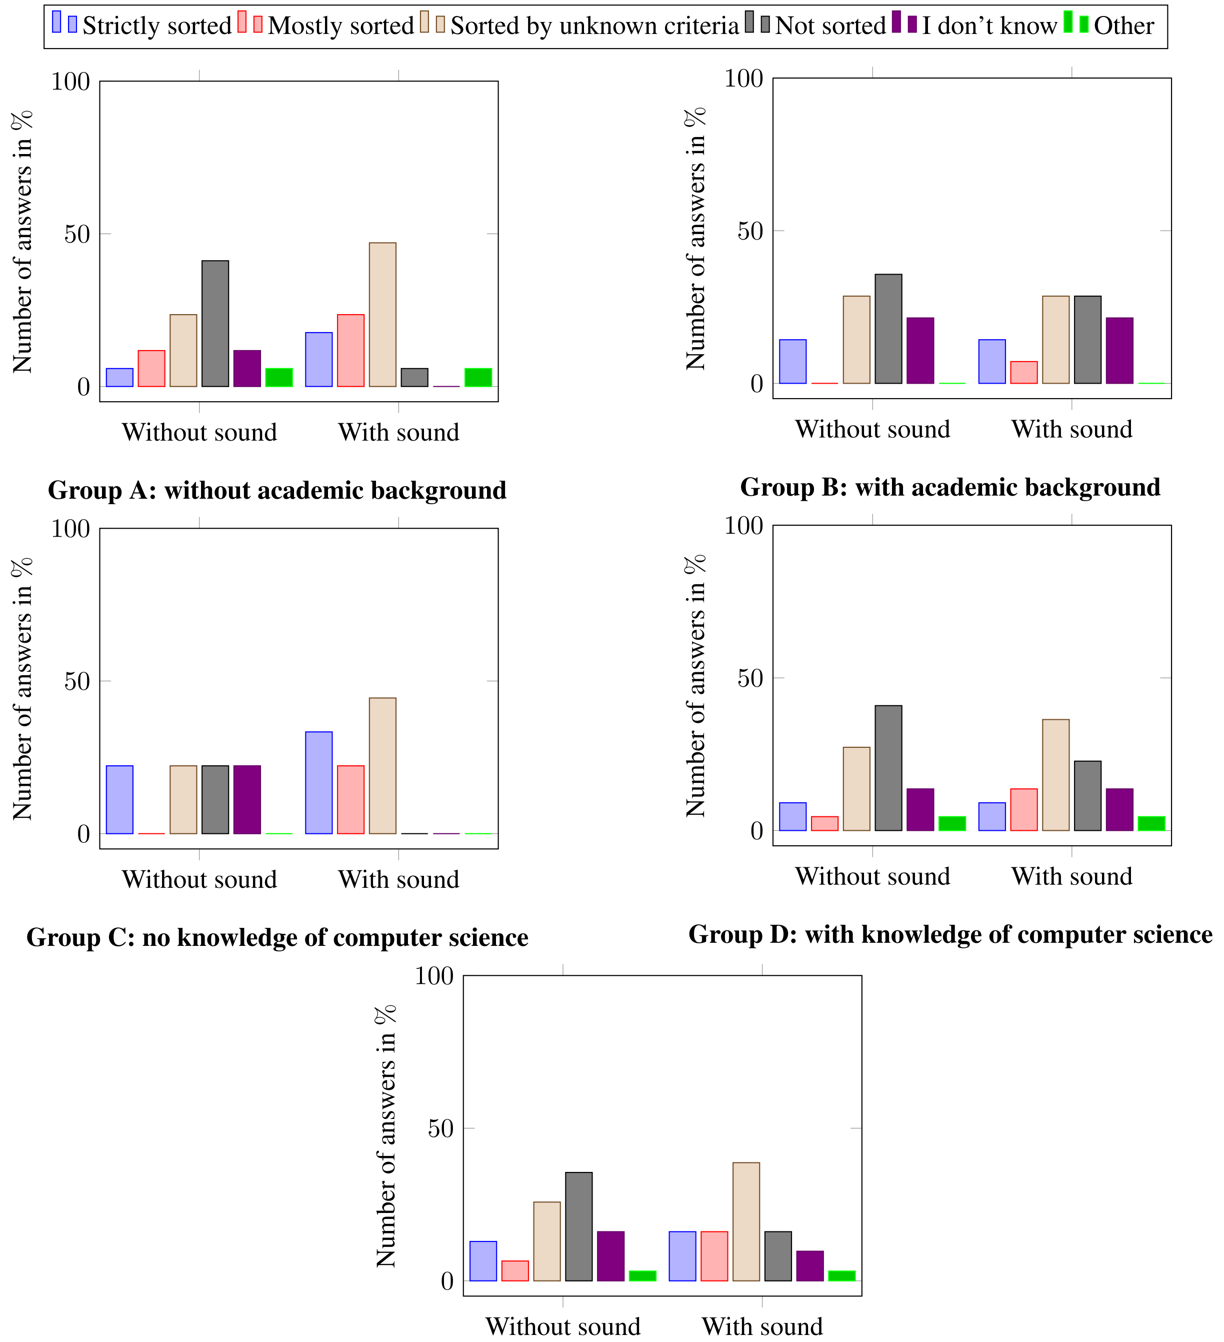

All groups

Figure 11: The results of the third task (sorting the data). The results are shown without sound on the left and with sound on the right.

recognized if the mapping is learned indirectly through the previous videos. The subjects are asked to use headphones to solve this problem. In the operator tree of the used query, the middle operator has twice as many tones as the "left" or "right" operators.

There was a very high increase in uncertainty in this task. The majority of all test persons answered *I don't know* with 38.7%. This answer option also dominates in groups B and D (B: 50%, C: 40.91%). In groups A and C, there was also a high level of uncertainty, but the answers for *3 operators, one particularly active* clearly predominate (A: 47.06%, B: 44.44%). In group

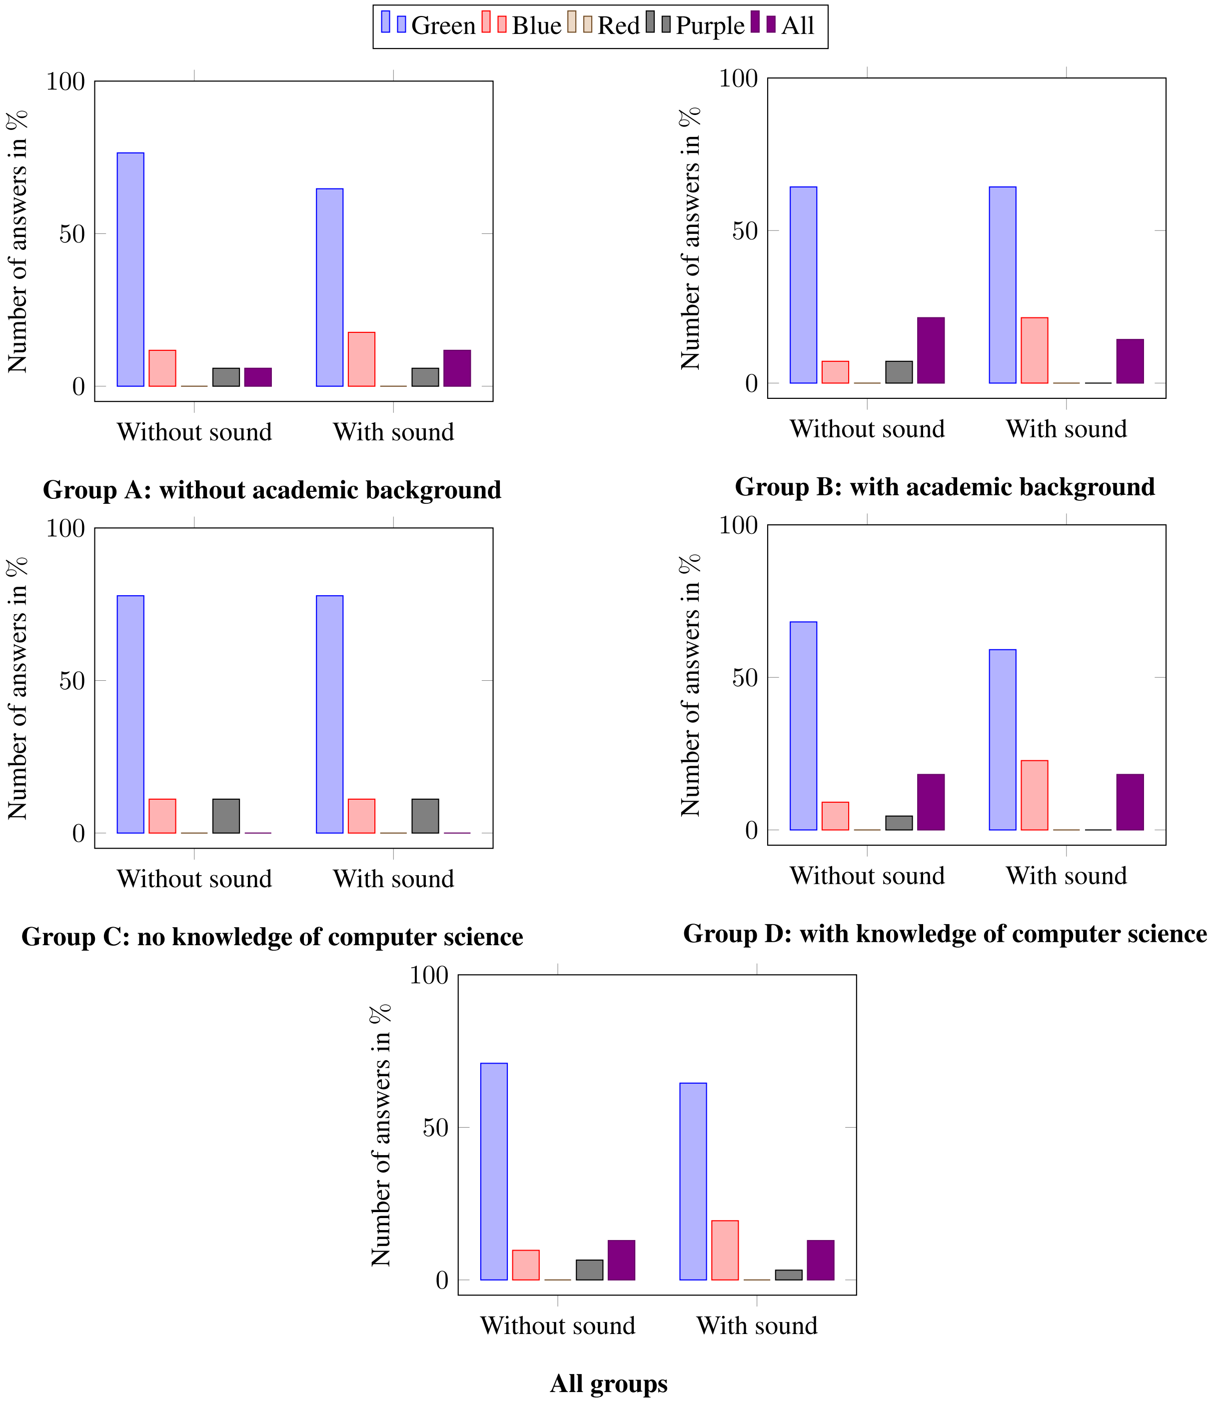

Figure 12: The results of the third task (recognizing different operators). The results are shown without sound on the left and with sound on the right.

B, compared to the other groups with 21.43%, most of the test persons stated that there are three operators, but that they are all equally active.

Overall, groups A and C performed significantly better than groups B and D in this task. Groups B and D were most uncertain. The high level of uncertainty can be explained by the fact that some of the test subjects failed to achieve the learning effect. However, it is also assumed that some test persons did not use headphones despite the recommendation and were thus unable to perceive the localization. The significantly better results in groups A and C can be explained by the fact that

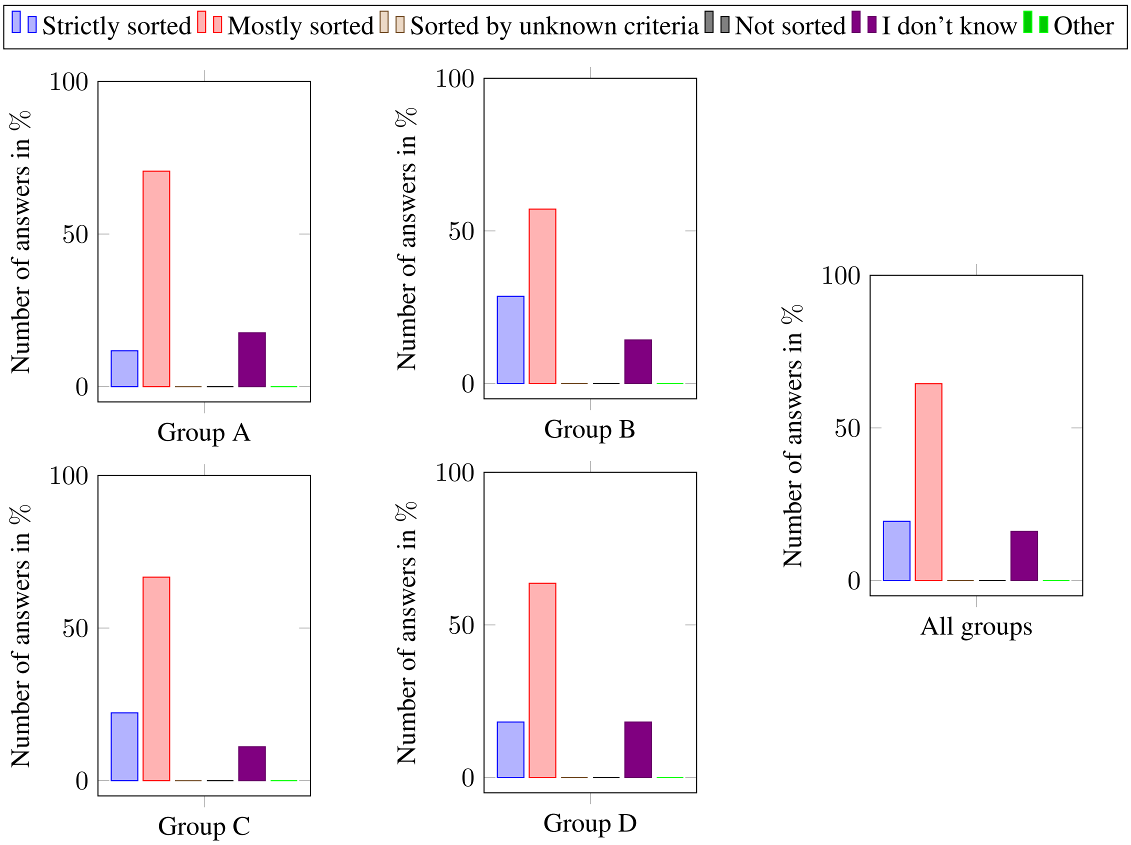

Figure 13: The results of the fourth task (sorting the data).

these groups are significantly more open to sonification. There is already a tendency for groups B and D to rely more on visual impressions. The test persons in groups A and C are not familiar with the subject of databases and operator trees, which is why pure visualization has less information content for these groups. Sonification therefore represents a greater added value than with groups B and D.

The results mentioned are presented in Figure 14.

Recognizing the join operators: By again mapping the guitar instrument to the join operator, the test subjects should recognize this special join operator. The guitar is used exclusively for the join operators in all of the videos shown. It can be assumed that this mapping would also cause a learning effect. As expected, the majority of the test persons answered the question about the existence of a join operator with *yes* (58.1%). However, these are significantly less than previously assumed, since the number of *I don't know* answers is significantly higher at 38.7%.

This answer does not occur as often in groups A and C (A: 29.41%, C: 33.33%) as in groups B and D (B: 50%, D: 40.91%), but is also significantly higher than expected. Since almost no test person opted for the answer option *No*, this high proportion of *I don't know* answers is again interpreted as uncertainty. The hoped-for indirect learning effect of the mapping turned out to be weaker than hoped, even if it can still be seen in the majority of the test subjects.

One reason for this uncertainty could be the choice of instrument for the join operator. Even if the piano and the guitar can be distinguished acoustically, the sound of the two instruments can be difficult to separate because they are somewhat similar in terms of the pitches used. With a different choice of instruments, the test subjects would probably have perceived the different instruments more strongly.

The answers of the test persons are shown in Figure 15.

Recognizing the sort operator: By mapping the xylophone instrument to the new sort operator unknown to the test persons, we want to check whether the test persons are able to recognize a new element in the operator tree, even if they have no visual information about it. The sort operator sorts the data values of the join operator from an ascending to a descending sort criterion.

Overall, the majority of the test persons succeeded in

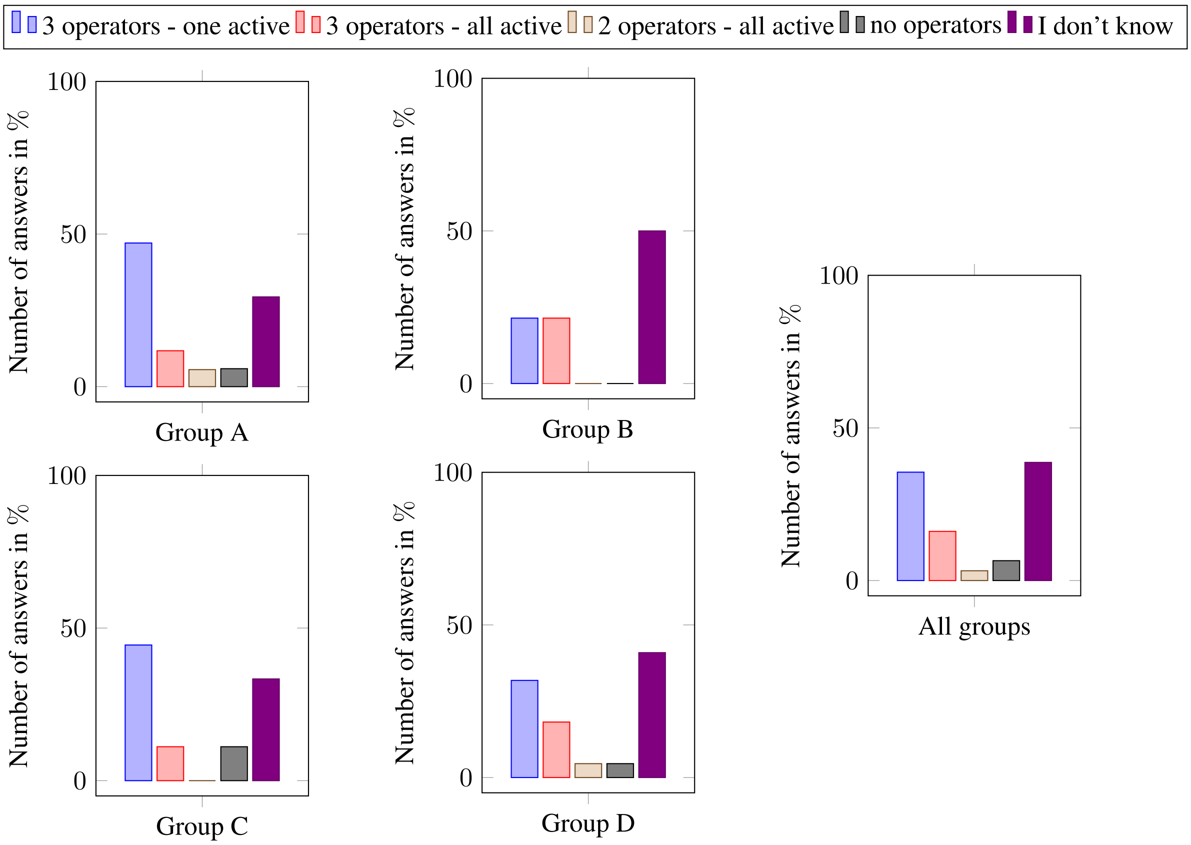

Figure 14: The results of the fourth task (recognizing different operators).

confirming the existence of an unknown operator with 64.5% (see Figure 16), even if there are big differences between the individual groups: In group B, the majority of the test persons answered *I don't know* with 50%. In group D this happens much less strongly with 31.82%, but this is a clear difference to groups A and C. Hardly any test person opted for this answer (A: 5.88%, C: 11.11%). Accordingly, the number of participants who chose the correct answer is also significantly higher in groups A and C than in groups B and D.

Recognizing the function of the sort operator: The results of the fourth task (recognizing the function of the sort operator) are shown in Figure 17). The sudden descending scale is intended to convey to the test persons that the unknown operator at hand is sorting the data in a different way. The sort operator generates as many data values as output as it receives as input. A filter function in which the operator only selects a subset of the data values can thus be excluded.

The correct answer was recognized by half of the test persons, while a quarter opted for the *filter function* and a further quarter said *I don't know*. Of all groups, group D indicated the *filter function* most with 33.33%. This assessment was lower in the other groups.

Group C was the most uncertain of all groups with 37.5%, while the uncertainty (*I don't know* answers) in the other groups was limited to $\leq 20\%$.

60% of the subjects from group B chose the correct answer option, with which group B performed best on this task.

The high level of uncertainty especially in group C are related to the fact that the test persons in this group are least familiar with filter and sorting functions in computer science. On the other hand, the two groups B and D in particular had less uncertainty and thus a greater number of correct answers, because both groups are much more familiar with this context.

8.2.5 Additional Questions

Finally, the test persons were asked to rate the individual statements A to G on a linear scale from 1 to 5, whereby 1 stands for "strongly disagree with the statement" and 5 for "strongly agree with the statement".

A) *The sonification of the query processing provides me a better overview:* This statement was answered very neutrally overall, whereby group B had the tendency that the sonification does not provide a

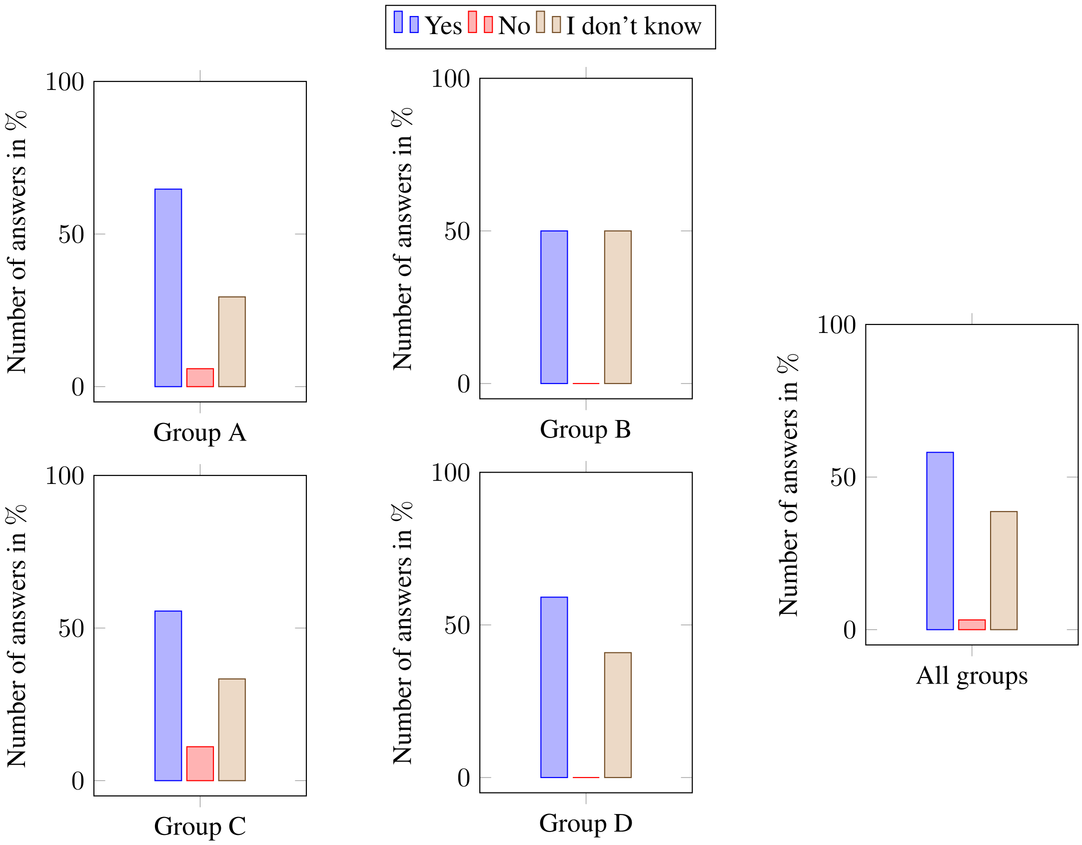

Figure 15: The results of the fourth task (join operator)

better overview. Groups A, C, and D indicate that the sonification tends to do this.

B) *The sonification of the query processing helped me to better answer the questions asked:* This statement was also evaluated neutrally. Group B again stated that the sonification was not helpful, while Group D rated this statement as neutral on average and Groups A and C rated this statement with a slightly positive tendency.

C) *The videos with the sonification are more interesting than the ones without the sonification:* For this statement all groups state that the videos with the sonification are more interesting than without the sonification, whereby groups B and D indicate this with a rating lower than groups A and C.

D) *Due to this survey I am more interested in databases:* This statement is again evaluated neutrally, with groups B and D indicating a negative tendency and groups A and C indicating a positive tendency.

E) *I find the idea of sonification of query processing exciting:* All groups regard the idea of sonification of query processing as exciting, whereby, similar to statement C, groups B and D rate this statement less positively.

F) *Due to the presented sonification I am able to remember better what the function of the join operator is:* With the rating to statement F the test persons should indicate whether the sonification ultimately helped better to keep the new information regarding the join operator. A positive tendency can only be seen in group C, while the other groups express themselves either neutrally or slightly negatively.

G) *I can imagine that artists could use this sonification to 'compose' new music:* With the last statement G almost all groups indicate with a slightly positive tendency that the proposed sonification can also be used for artistic purposes. Only group B rated this statement with a negative tendency.

The evaluations are presented in Figure 18.

A tendency can be observed across all statements: Group B basically rated all statements with the fewest points, closely followed by group D. The ratings of groups A and C are always higher.

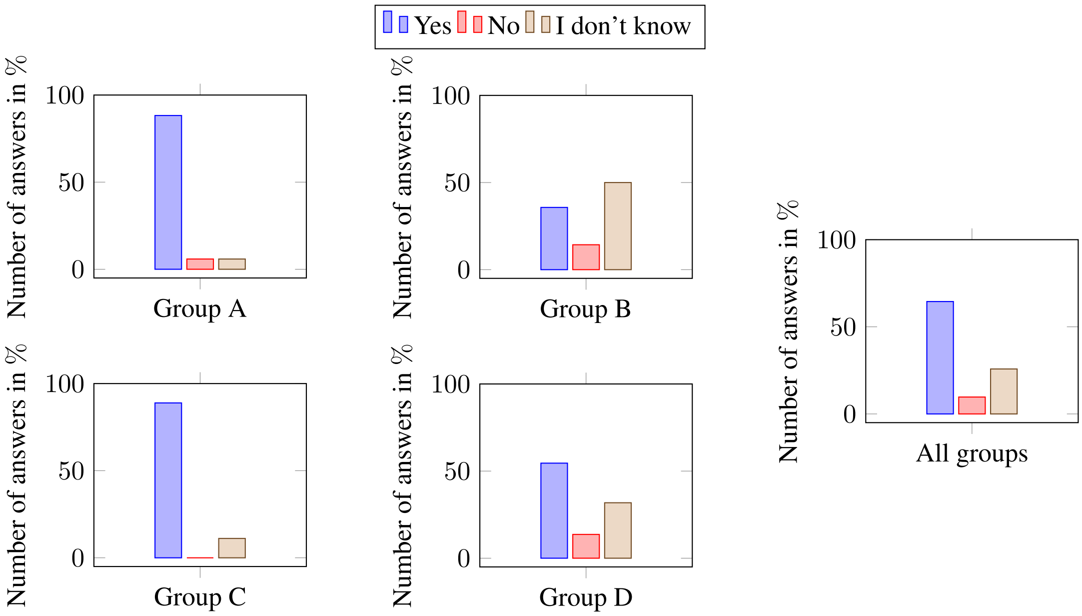

Figure 16: The results of the fourth task (unknown operator)

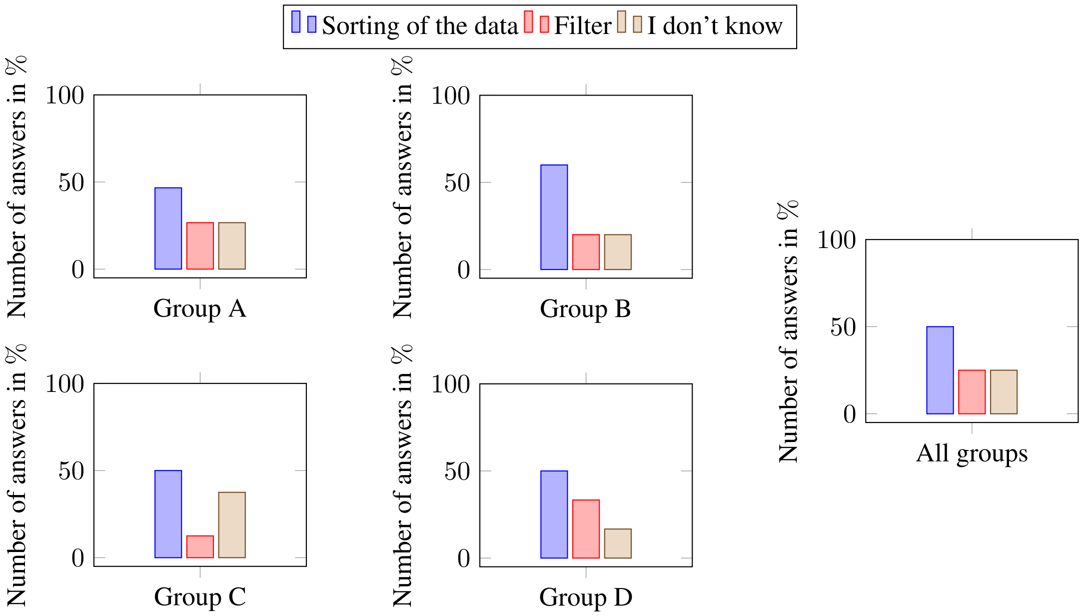

Figure 17: The results of the fourth task (function of new operator).

8.3 Discussion of the Results of the Online Survey

In the following, the hypotheses set up in section 8 will be validated and discussed on the basis of the acquired results of the online survey.

Hypothesis A: Sonification has a positive effect on the analysis of database data and pattern recognition. First of all, it must be mentioned that the test persons were not informed of any of the selected mappings of operator tree information to auditory dimensions. These mappings should be learned indirectly. If you look at the results of the sorting tasks for tasks 1 to 3, you can see that the results have always improved with the sonification. For example, 51.5% stated in task 1 that the data was not sorted at all, while the sonification led to an improvement, so that only 16.1% voted for non-sorting. Similar results can also be seen in the sorting

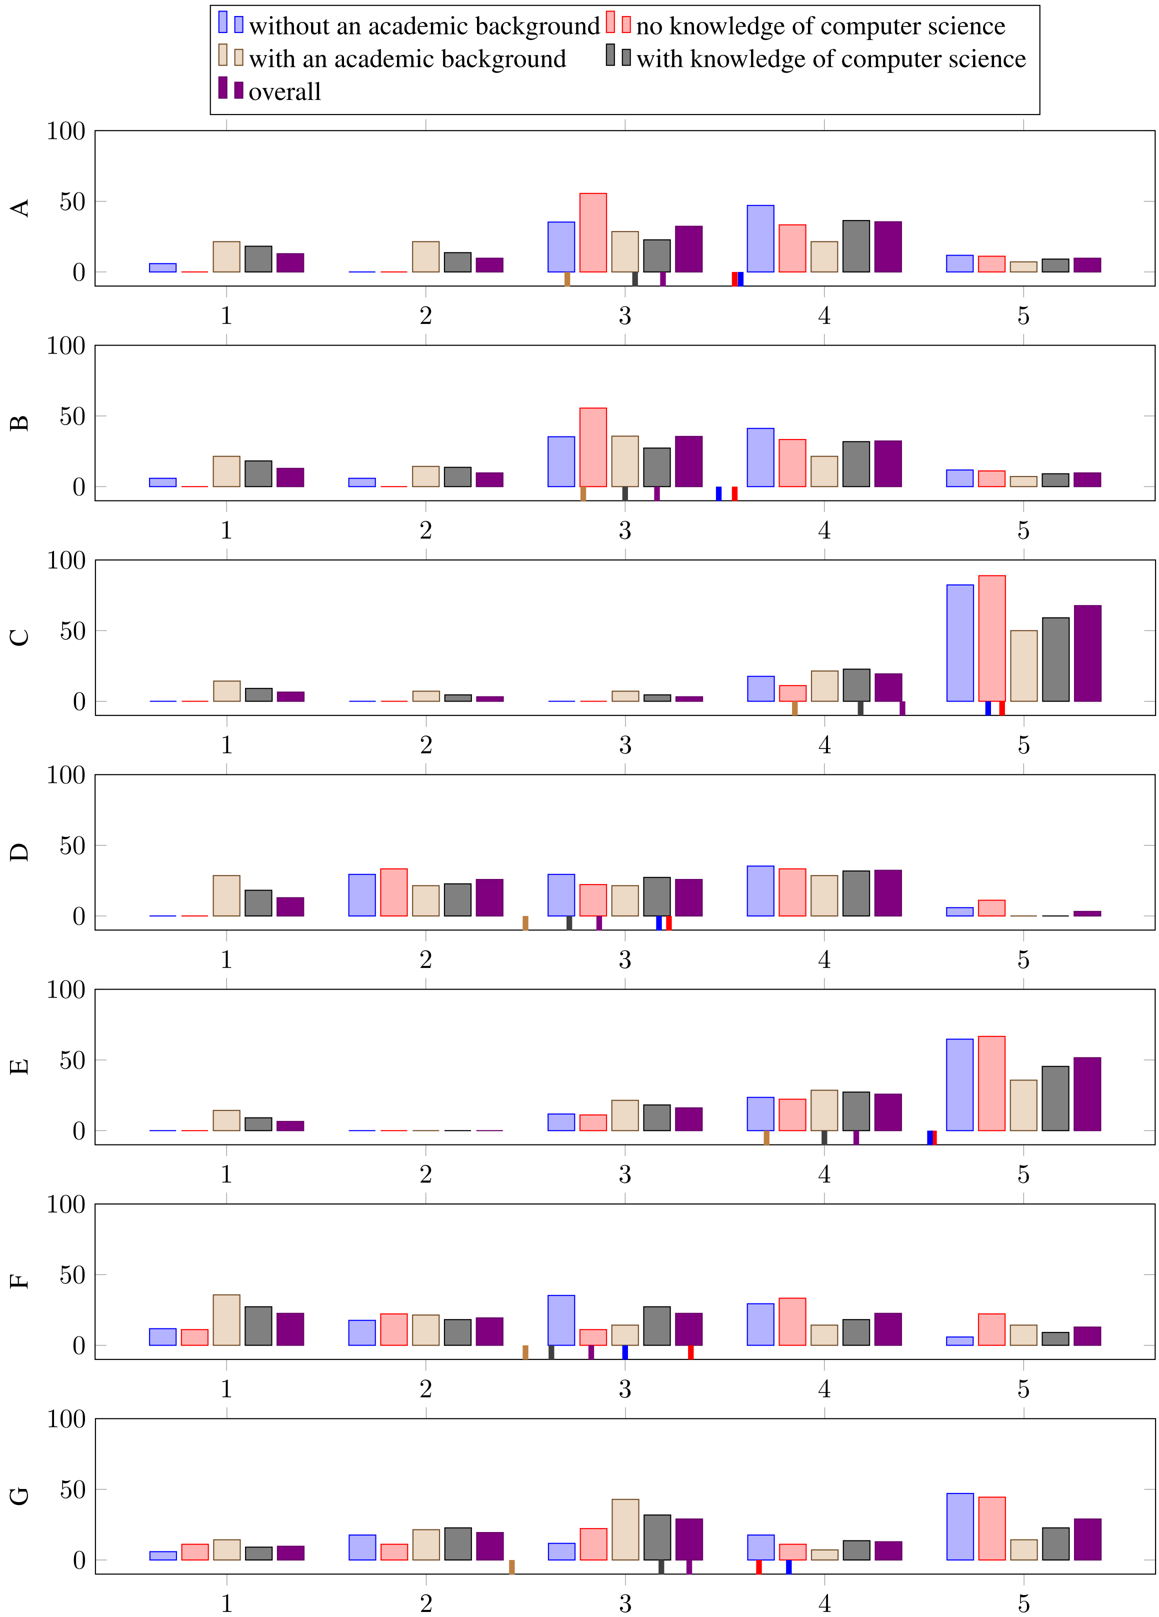

Figure 18: The evaluations of the different groups for the statements A-G. The mean values are marked on the x-axis.

tasks of tasks 2 and 3: In task 2, the majority of the test persons corrected their statement from a strict sorting to the correct answers, while the number of answers from *not sorted* decreased. This can be seen particularly clearly in exercise 3, with the answers *not sorted* change from 35.5% to 16.1%, while the correct answer option was chosen significantly more often.

The purpose of task 4 is to determine the extent to which the selected mappings are learned unconsciously and the extent to which the subjects are able to solve the tasks even without visual information. Most of the test persons were able to choose the correct answer for the sorting task with 83.9%, although they had to rely purely on their hearing. Even the task of recognizing the operators, the mapping of which is not so obvious, could still be solved correctly with at least 35.5%, where only the answers with *I don't know* predominate. Around 64.5% recognized by the choice of the instrument that a join operator was available again and 64.5% also recognized a completely unknown operator, of which 50% could only determine the function of the operator through the acoustic information.

The uncertainty has increased significantly in the 4th task in particular. The test groups that are used to dealing with scientific data also tend to be more confused, which was particularly evident in the fourth task. It can be assumed, however, that all test persons would have been significantly more confident in answering the questions if they had been made aware of the mappings and thus knew which properties of the sonification to pay attention to.

The typical user of the LUPOSDATE3000 SPA client sets their mappings to auditory dimensions themselves and adapts them to their preferences and needs. Therefore, the user knows exactly which acoustic elements represent which information. It can therefore be assumed that the majority of the people who were unsure in this online survey would also see a benefit in the sonification in real operation.

However, the sonification alone does not seem to be particularly helpful when it comes to solving unknown, complex issues. Anomalies such as deviating sorting patterns or the sudden appearance of a new instrument can be recognized, but the answers to the question about the specific sorting were widely scattered.

Provided that the underlying mappings to auditory dimensions are known and the user is allowed to experience a certain amount of training, a combination of visualization and sonification seems to be suitable for performing analysis tasks and pattern recognition. The online survey was thus able to confirm hypothesis A.

Hypothesis B: Sonification helps to make scientific data more understandable for non-academic people. For the validation of this hypothesis, we have to primarily focus on group A's answers. This group stated that the sonification help to create a better overview and make it easier to answer the questions. However, this assessment is only slightly positive. The sonification does not increase the interest in databases themselves, but the sonification makes the whole topic much more interesting and exciting for this group. In fact, both groups A and C are able to get better results than groups B and D on most of the tasks.

In general, the visualization, together with the sonification, opens up a whole new level with which people can interact with the system and understand the processed operations. The willingness to deal with query processing is significantly higher than the conventional display based purely on text and tables.

The positive evaluations of statements A, B, C and D by groups A and C as well as the consistently positive results show that the sonification of data processing can significantly help to make scientific data accessible to a large part of the population who are or are not familiar with the respective scientific context. Of course the way of the presentation as well as the interaction with the data should take place in an appropriate and understandable framework.

In summary, hypothesis B is confirmed if it is ensured that the sonification and the associated overall context is prepared in a way that is understandable for the general public.

Hypothesis C: Sonification is also suitable in the context of databases for visually impaired people. In this online survey, there was no experimental group representing visually impaired people. Accordingly, there is no data basis to make a general statement here. Task 4 aims specifically at answering this hypothesis, which is why the test subjects are deprived of the visual channel for answering.

In task 4, the test persons are only able to achieve good and correct results by using the acoustic channel. The increasing uncertainty in the 4th task is due to the lack of special training. If the test subjects are given enough time to familiarize themselves with the system, far better results can be expected. The fact that the subjects in the 4th task have the same prerequisites as people with a visual impairment suggests that a test group with visually impaired people would achieve similar results.

It must be noted, however, that the questions refer to relatively simple facts. In practice, analysis tasks are far more complex than the general determination of a sorting or the recognition of an operator. The sonification as implemented in this contribution would therefore only be able to help visually impaired people to a limited extent to work with databases in practical situations. However, it can be assumed that a system can be helpful if it is appropriately adapted to the special needs of visually impaired people.

With the proviso that more work has to be invested in the development for a really helpful application, hypothesis C can also be confirmed in the context of databases.

Hypothesis D: Sonification has a positive learning effect for pupils at school and students. The sonification lead to better results for almost all tasks than for the purely visual tasks. This fact alone suggests that correspondingly better learning successes can be expected for pupils at school and students if the sonification of the processes in a database is integrated into the teaching process. Here, too, a corresponding training period is required in order to familiarize the pupils and the students with the sonification.

The evaluations of statement F show that the sonification alone is not enough so that complex content can be learned. Only group A reports a slightly positive learning effect about learning how the join operator works.

If sonification is only used as an additional medium that supports conventional methods, but is not intended to replace them, then hypothesis D can also be confirmed.

Hypothesis E: Sonification can be used to create artistic works from data. As is already known from the literature, sonification in general has already been used for some artistic productions. The online survey, on the other hand, is intended to examine whether this is also possible in the context of databases.

If the test persons are asked explicitly for their assessment of this statement, groups A and C give a rating with a positive tendency. Group D remains neutral in the evaluation and group B evaluates this question with a negative tendency. However, all groups state that the proposed sonification increases their interest in the related topics and that sonification of data processing is exciting.

From the direct results of the online survey using the LUPOSDATE3000 SPA client hypothesis E cannot be directly confirmed. A corresponding preparation of the representation and the elaboration of a particularly interesting mapping to sound synthesis parameters can, however, have the potential to lead to an artistic staging. Whereas we missed to play some good examples of music sounds generated by the LUPOSDATE3000 SPA client in the online survey, we played some of these examples[1] at the VLDB conference in 2021 and got very positive feedback.

Under the above conditions, hypothesis E can also be confirmed.

9 Summary and Conclusions

In this contribution we deal with the sonification of query processing. We showcase the sonification of numerous queries by proposing a sophisticated web application supporting a simple way of a complex configuration of the mapping from query processing steps to auditory dimensions. With the help of the sonification, we offer a unique audible experience and new way of experiencing query processing. The listeners of our sonification are encouraged to guess the functionality of the different operators in the queries and determine the types of processed relational operators. Furthermore, we show that auditive analysis are possible by detecting common subexpressions and similar queries.

We analyze the effects of sonification in an extensive user evaluation and analyze the results of an online survey. The analysis confirms that when sonification of data processing in databases is offered as additional medium, it has a positive effect on the analysis of database data and pattern recognition, helps to make scientific data more understandable for non-academic people, has a positive learning effect of database algorithms and can be used to create artistic works from data.

In future work, we will consider other open source databases and database tasks in the area of parallel of distributed processing as well for integration into our sonification approach.

Acknowledgements

This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 422053062.

References

[1] G. Aad, J. Butterworth, J. Thion, U. Bratzler, P. Ratoff, R. Nickerson, J. Seixas, I. Grabowska-Bold, F. Meisel, S. Lokwitz *et al.*, "The atlas experiment at the cern large hadron collider," *Jinst*, vol. 3, 2008.

[2] S. Barrass, "Sonification design patterns." Georgia Institute of Technology, 2003.

[3] L. Commère, G. Florian, F. Côté, and J. Rouat, "Sonification of 3d point clouds for substitution of vision by audition for blind users," in *24ème conference internationale Auditory Display (ICAD 2018)*, 2018.

[4] G. Dubus and R. Bresin, "A systematic review of mapping strategies for the sonification of physical

quantities," *PloS one*, vol. 8, no. 12, p. e82491, 2013.

[5] J. H. Flowers, D. C. Buhman, and K. D. Turnage, "Data sonification from the desktop: Should sound be part of standard data analysis software?" *ACM Trans. Appl. Percept.*, vol. 2, no. 4, p. 467–472, 2005. [Online]. Available: https://doi.org/10.1145/1101530.1101544

[6] S. Groppe, *Data Management and Query Processing in Semantic Web Databases*. Springer, 2011. [Online]. Available: https://doi.org/10.1007/978-3-642-19357-6

[7] S. Groppe, "Semantic hybrid multi-model multi-platform (shm3p) databases," in *Proceedings of the International Semantic Intelligence Conference (ISIC), New Delhi, India*, 2021. [Online]. Available: http://ceur-ws.org/Vol-2786/Paper2.pdf

[8] S. Groppe and J. Groppe, "Hybrid multi-model multi-platform (hm3p) databases," in *Proceedings of the 9th International Conference on Data Science, Technology and Applications (DATA), virtual due to Covid-19 pandemic*, 2020. [Online]. Available: https://doi.org/10.5220/0009802401770184

[9] S. Groppe, R. Klinckenberg, and B. Warnke, "Sound of databases: Sonification of a semantic web database engine," *Proc. VLDB Endow.*, vol. 14, no. 12, pp. 2695–2698, 2021. [Online]. Available: https://doi.org/10.14778/3476311.3476322

[10] O. Hans-Martin Lutz, J. L. Kroger, M. Schneiderbauer, and M. Hauswirth, "Surfing in sound: Sonification of hidden web tracking." Georgia Institute of Technology, 2019.

[11] B. Hansen, L. B. Baltaxe-Admony, S. Kurniawan, and A. G. Forbes, "Exploring sonic parameter mapping for network data structures." Georgia Institute of Technology, 2019.

[12] T. Hermann, A. Hunt, and J. G. Neuhoff, *The sonification handbook*. Logos Verlag Berlin, 2011.

[13] E. Hill, J. Cherston, S. Goldfarb, and J. Paradiso, "Atlas data sonification: a new interface for musical expression," in *38th International Conference on High Energy Physics*, vol. 282. SISSA Medialab, 2017, p. 1042.

[14] W. Joo, "Sonifyd: A graphical approach for sound synthesis and synesthetic visual expression." Georgia Institute of Technology, 2019.

[15] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. Flowers, N. Miner, J. Neuhoff, R. Bargar, S. Barrass, J. Berger *et al.*, "The sonification report: Status of the field and research agenda. report prepared for the national science foundation by members of the international community for auditory display," *International Community for Auditory Display (ICAD), Santa Fe, NM*, 1999.

[16] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. Flowers, N. Miner, J. Neuhoff, R. Bargar, S. Barrass, J. Berger, G. Evreinov, W. T. Fitch, M. Gröhn, S. Handel, H. Kaper, H. Levkowitz, S. Lodha, B. Shinn-Cunningham, M. Simoni, and S. Tipei. (1997) Sonification report: Status of the field and research agenda. [Online]. Available: http://www.icad.org/websiteV2.0/References/nsf.html

[17] B. Laugwitz, T. Held, and M. Schrepp, "Construction and evaluation of a user experience questionnaire," in *Symposium of the Austrian HCI and usability engineering group*. Springer, 2008, pp. 63–76.

[18] D. E. Millard and M. Ross, "The 3d sonification of links in physical hypermedia environments," in *Proceedings of the sixteenth ACM conference on Hypertext and hypermedia*, 2005, pp. 251–253.

[19] J. G. Neuhoff, "Is sonification doomed to fail," in *The 25th International Conference on Auditory Display (ICAD), Newcastle-upon-Tyne, UK*, 2019.

[20] T. Neumann and G. Weikum, "The RDF-3X engine for scalable management of RDF data," *VLDB J.*, vol. 19, no. 1, pp. 91–113, 2010. [Online]. Available: https://doi.org/10.1007/s00778-009-0165-y

[21] M. Porter. (2020) Data sonification: A new cosmic triad of sound. [Online]. Available: https://www.nasa.gov/mission_pages/chandra/news/data-sonification-a-new-cosmic-triad-of-sound.html

[22] M. Quinton, I. McGregor, and D. Benyon, "Investigating effective methods of designing sonifications." Georgia Institute of Technology, 2018.

[23] A. G. Riber, "Sonifigrapher. sonified light curve synthesizer," 2019.

[24] N. Rönnberg and J. Johansson, "Interactive sonification for visual dense data displays," in *Proc. 5th Interactive Sonification Workshop (ISON-2016)*, 2016, pp. 63–67.

[25] N. Sawe, C. Chafe, and J. Treviño, "Using data sonification to overcome science literacy, numeracy, and visualization barriers in science communication," *Frontiers in Communication*, vol. 5, p. 46, 2020.

[26] M. Sosby. (2019) Data and music: What 50 years of exploring our moon sounds like. [Online]. Available: https://www.nasa.gov/feature/goddard/2019/data-and-music-what-50-years-of-exploring-our-moon-sounds-like

[27] K. Vines, C. Hughes, L. Alexander, C. Calvert, C. Colwell, H. Holmes, C. Kotecki, K. Parks, and V. Pearson, “Sonification of numerical data for education,” *Open Learning: The Journal of Open, Distance and e-Learning*, vol. 34, no. 1, pp. 19–39, 2019.

[28] B. Warnke, M. W. Rehan, S. Fischer, and S. Groppe, “Flexible data partitioning schemes for parallel merge joins in semantic web queries,” in *Datenbanksysteme für Business, Technologie und Web (BTW), Dresden, Germany*, ser. LNI, 2021.

AUTHOR BIOGRAPHIES

Sven Groppe earned his diploma degree in Informatik (Computer Science) in 2002 and his Doctor degree in 2005 from the University of Paderborn. He earned his habilitation degree in 2011 from the University of Lübeck. He worked in the European projects B2B-ECOM, MEMPHIS, ASG and TripCom. He was a member of the DAWG W3C Working Group, which developed SPARQL. He was the project leader of the DFG project LUPOSDATE, an open-source Semantic Web database, and one of the project leaders of two research projects, which research on FPGA acceleration of relational and Semantic Web databases. He is also leading a DFG project in the Semantic IoT area and a DFG project on GPU and APU acceleration of main-memory database indexes. He is also the chair of the Semantic Big Data (SBD) workshop series (2016 to 2020) and Big Data in Emergent Distributed Environments (BiDEDE) in 2021, both are affiliated with the ACM SIGMOD conference, and of the Very Large Internet of Things (VLIoT) workshop in conjunction with the VLDB conference (so far 2017 to 2021). His research interests include databases, Semantic Web, query and rule processing and optimization, Cloud Computing, acceleration via GPUs, FPGAs and quantum computers, peer-to-peer (P2P) networks, Internet of Things, data visualization and visual query languages.

Rico Klinckenberg was born in Ribnitz-Damgarten, Germany in 1995. He received his M.Sc. in Entrepreneurship in Digital Technologies in 2021 from the University of Lübeck, Germany. He wrote his master thesis at the institute of information systems (IFIS) about the sonification of databases where he mainly focused on developing a new sonification tool, which maps the visual information of the operator tree of a triple store to audio. He used this application to investigate which benefits the sonification can have in the context of databases.

Benjamin Warnke has received his M.Sc. in Computer Science in 2019 from the University of Hamburg, Germany, with a thesis about about integrating self-describing dataformats into file systems. At the moment he is employed as a research assistant at the Institute of Information Systems at the University of Lübeck, where he is working on the Semantic Web Database Luposdate3000. His research interests include space and time efficient Data Structures, databases and file systems.