TOWARDS AN AGENDA FOR AUDITORY OVERVIEWS

Tony Stockman and Louise Valgerður Nickerson

Cognitive Science Research Group
Queen Mary University of London
Mile End Road, London
{tonys,lou}@eecs.qmul.ac.uk

ABSTRACT

We examine the use of overviews in the context of auditory displays. We begin by addressing what overviews are and what is their role within the broader context of Human-Computer Interaction. This leads to the identification of a set of characteristics which, based on a literature survey and our own analysis of commonly occurring overviews we find that overviews possess. We then examine to what extent these characteristics are present in several reported examples of auditory overviews. It is noted how much this analysis could be improved in a research environment which fostered repeatability and comparability. The paper concludes by arguing why overviews are particularly valuable in auditory displays within the increasingly important contexts of *Big Data* and mobile use.

1. INTRODUCTION

This paper argues that auditory overviews are an under exploited mechanism within the field of auditory displays, and that they have a potentially extremely important part to play within the context of *Big Data* and increasingly intelligent mobile devices with small screens. We start by reviewing the nature of overviews in general, and by reviewing some notable examples of overviews specifically designed for auditory displays. We analyse these examples to establish what features they share with some widely known and understood visual overviews. Finally we outline an agenda for the future development of auditory overviews, setting out a number of areas where we believe they have an important contribution to make.

2. WHAT IS AN OVERVIEW?

Shneiderman and Plaisant have written widely on the field of overviews in the context of information seeking. They and their colleagues [1, 2, 3, 4, 5, 6] have used overviews for a variety of applications, resulting in Shneiderman's seminal paper citing overviews [3] as an integral part of the interface and as the first step in exploring a data set in the information seeking mantra ("overview first, zoom and filter, and then details-on-demand"). [1] compares an overview and two other interfaces for browsing hierarchies. In [7], overviews are applied to personal histories to help highlight connections between otherwise disparate events. [4] describes a novel overview for photo libraries. In all these cases, overviews are shown to benefit the interface, however the topic of what an overview is and what role it should play is not addressed directly. [8] adapts Shneiderman's approach to information seeking to auditory interfaces, substituting the word *gist* for the word *overview* to avoid visual nomenclature. Similarly, [9] adapted the information seeking approach of Shneiderman to an auditory context, this time substituting the word *overview* with *situate*. The situate command informs users of the structure of the page and available options based on the location of the cursor. [10] provided auditory overviews of program source code where the overview describes the hierarchical structure of nested statements in the code. [11] described an overview of node/edge diagrams which showed the size and complexity of the graph and highlighted the interrelationships between nodes. [12] focused on the auditory representation of numerical tabular data. The overview of rows and columns allowed users to identify areas of interest to explore. While the term overview is used in all of these, it is taken for granted that the reader knows what an overview is and how it can benefit the interfaces described.

2.1. Attributes of overviews from the literature

The literature reviewed above deals with overviews, however it deals with them without delineating what they are. The purpose of this review is to draw out what can be generalised about overviews.

Comprehensive: Overviews describe an entire collection [3, 7, 10, 11] of information. An exception is however described in [12] where the user directs how smaller overviews can be put together to form an overview of the entire collection, so leading to the concept of a hierarchy of overviews.

Abstraction: Overviews provide a general understanding of the detailed data, obscuring detail and reducing complexity. [7, 10, 12, 13]

Guide the user: Overviews are important for navigation and point the user in the right direction to find what they are looking for. [7, 9, 13]

Displays saliency/interrelationships: Overviews expose how the detailed data is interconnected, give details a frame of reference and identify areas of interest. [7, 11, 12, 13]

Other characteristics that emerged less frequently are that: overviews promote exploration [13] by preventing users from getting lost in the data; they expose the structure of the detailed data [9]; and they are separate from the detail [14].

2.2. Survey of overviews

To examine the validity of the above analysis, we turn to a more in depth analysis of overviews themselves. We analyzed known overviews in order to identify some of their common attributes. We chose four common overviews in a variety of formats so that we

could properly generalize our findings. The overviews are: tables of contents, computer file managers, line graphs and abstracts. The first three can be automatically generated while abstracts cannot; three are textual while line graphs are graphical; and all are well-understood so that the analysis is uncomplicated by their purpose.

2.2.1. Tables of contents

Tables of contents are textual and generally automatically generated from the structure of the document: headers are extracted from the text and displayed along with their associated section and page numbers. Components and their representation Copy-editing: The Cambridge Handbook for Editors, Authors and Publishers [15] indicates that tables of contents should be comprehensive but simple to read. Some of the typical components are shown in Table 1.

Component	**Role**
Headings/sub-headings	shows topics covered
Sets of headers	flow of topics
Heading numbers	structure of the material
Page numbers	where material is located also depth in which material is covered
Indentation	differentiates major topics from minor topics
Leaders	legibility: which pages numbers correspond to what topic

Table 1: **Components of a table of contents and their roles in the overview**

Superficially, tables of contents look dry and simple: they are listings of the top level headers of a document. However, this overview goes beyond an enumeration. Part of this is visible from the interaction between the overview and the detail during creation of the document [16]. The table of contents (or outline) can affect the writing of a document since the document itself would need to change in order for the overview to better represent it. That the overview shows the author a different viewpoint on her/his document is mirrored in the reader's experience. The reader can discover things that s/he did not know s/he was looking for (similar findings are described in [17]). In both instances, the overview encapsulates what the document is about and acts as a guide. The reader sees what topics surround a heading (or topic) and this gives her/him a sense of how the topic is handled, in what depth and engages her/him to explore what else is contained in the document. Page numbers, indentation and other formatting all convey minute details about the content of the document that a list cannot. It is the organisation of the information that makes the overview useful; otherwise, a search feature would be sufficient for exploring the document. This concept of the table of contents adding to the reader experience is discussed in [18].

Based on this analysis of tables of contents, overviews include the following qualities: setting out the scope and structure of the material, delineating the hierarchy of the material (if one is present), quickly guiding the reader to topics of interest, showing contextual information about topics, helping the reader find where s/he left off in prior interaction and encouraging the reader to explore other topics.

2.2.2. File managers

File managers employ graphical elements to represent the organisation of a computer's file system. The file manager in the Windows[1] operating system displays the current path in the address bar, the contents of that directory in the main panel and an interactive representation of the directory tree structure on the left. This representation of the directory tree structure is an overview that describes the entire contents of the file system where the user can choose which portions to hide or expand depending on her/his task. The main advantage of the overview is the visibility of the entire file system where previously the user would have to stitch together her/his own internal representation of the surrounding context.

Components and their representation The strength of file managers comes from the way they display the relationships between folders so that users are familiar with the organisation and have an easier time remembering or deducing where they stored certain files. the directory tree is mostly made apparent through the spatial layout that indicates the relationships between folders. The components of the file manager are shown in Table 2

Component	**Role**
Folders and sub-folders	shows folders in the directory tree
Sets of parent/children folders	organisation of information
Sibling folders	folders loosely associated with each other
Parent folders	general category of a folder
Expanded icon	state of the folder (expanded or hidden)
Indentation	folder depth in the directory tree

Table 2: **Components of a file manager and their roles in the overview.**

The directory tree overview is based on the metaphor of the real-world desktop and office environment: The file manager is a filing cabinet, directories are folders and the contents are files. This is a much newer overview than tables of contents and has received much criticism (e.g. [19, 20, 21]). The main criticism is the difficulty of finding files efficiently.

This criticism indicates the immaturity of the overview and how it could benefit from improvement. However, the criticism is perhaps more directed at the organisation of the underlying files rather than the abstraction thereof. One flaw in the overview itself is a break in the hierarchical organisation: the Windows file manager lists the Desktop folder twice. It appears as the top level directory as well as a child of the user's home directory. Here the overview misrepresents the underlying structure and potentially causes confusion. This is handled differently in OSX[2], where the main overview is divided into sections. It provides shortcuts or hooks into the file system rather than displaying its organisation. However, as we discuss below, this view does not describe the file system fully, which seems to be an important part of an overview.

Another problem with the overview in Windows is, as sections of the file system are expanded, portions of the directory tree structure disappear off the edges of the overview pane. Having to scroll

[1] http://www.microsoft.com
[2] http://www.apple.com/macosx/

can result in confusion until the user finds her/his place again in the view. However, this is more of an implementation usability problem – where there is a trade-off in usage of available screen space – than a problem with the overview per se.

As mentioned above, the Apple file manager takes a different approach. The default view uses a series of panes where each pane lists a level in the directory tree but never displays the entire file system structure. This leaves the question of whether this representation qualifies as an overview since it acts as a filter. Per Shneiderman's info-seeking mantra filtering is a latter step in the process of exploration. More specifically, it acts more as a focus+context interface, where the user sees the current detail in full focus but also gets contextual information. [22] The distinction is that the context provides a view of surrounding information but not of the whole.

So what does the Windows file manager imply about overviews? The analysis of tables of contents revealed that an overview sets out scope, structure and hierarchy. The file manager does these as well though the interaction degenerates when parts of the overview scroll out of its pane. The exploratory aspects are also less evident. One new attribute that emerges is the affect of the overview on a dynamic system. File systems are in constant flux and as such snapshots of the overview can describe the evolution of the contents of the file system.

2.2.3. Line graphs

Line graphs show interrelations between two measures. Their usage in statistics and analysis made them a key feature in spreadsheet programs such as Excel[3] and they are well enough understood to be able to be generated automatically.

Components and their representation The main component of a line graph is the line which is made up of connected data points. The positions of the data points are determined by the scales of the two axes. The axes will have labels and units of measurements so that one can tell what the data points mean. More advanced line graphs might have trend lines and/or error bars to assist interpretation. The main components – but not these more advanced ones – are in Table 3

Component	**Role**
Line	connect the data points and show the progression or variation of the data along the two axes
Data points	show the specific values for each measurement
Axes	show what is being compared
Axis labels and units	describe the axes and their scale
Grid lines/tick marks	show the scale of the axes and help people to approximate data point values
Title	describe the subject of the graph
Legend	labels the data line in the graph

Table 3: **Components of a line graph and their roles in the overview.**

In terms of the previously exposed attributes of overviews, line graphs describe the scope and structure of the data. In addition, they quickly guide and entice the viewer and provide a snapshot of the detailed information. However, they sometimes only describe a subsection of an entire data set. When the data set is small, a graph can describe it in its entirety but as the data increases in dimensions, it may only describe one aspect. Regardless of the scope of the line graph, it can depict a multitude of information. Like a table of contents, a line graph invites exploration. The study of the graph can lead to questions and subsequent searches for answers. Graphs can both show findings and identify areas of interest to analyse. [23] discuss graphs as tools in exploratory data analysis, specifically highlighting their ability to provide "a good view of the relationships and oddities in the data from experiments" [p. 120]. As such, they have the potential to expose aspects of the detailed information that are not immediately obvious without the benefit of an overview.

[3] http://office.microsoft.com/en-us/excel/

2.2.4. Abstracts

Abstracts are textual, like tables of contents, yet are hand-generated. This is in contrast to the previous examples which can all be generated automatically from the detailed information. Components and their representation Abstracts are very specialized and good ones are carefully constructed to best represent the work in question. In the most general terms, an abstract will have an introductory statement, a main message and an explanatory section. The order and format depends greatly on the type of work, the forum/audience and the perspective of the author.

Component	**Role**
Introductory statement	describes the purpose of the work and sets the scene
Main message	what the author wishes the reader to remember
Explanatory section	describes background information and/or methodology

Table 4: **Components of an abstract and their roles in the overview.**

Because of the prose format, abstracts are hard to create and hard to describe. They are short and the author must be very concise. Unlike the other overviews studied, abstracts resist automatic generation. As such, abstracts are a good example of how one needs understanding of the underlying work to be able to generate a new overview. The way we express ourselves in words is not yet well enough understood to negate the need for a custom overview. Similarly, in cases where we are aiming to form new overviews, overviews will need to be custom generated until the underlying format is well enough understood to be automatically processed.

3. APPLYING THE RESULTS TO THE LITERATURE

The survey reported above identified several key characteristics of overviews. Not all of them appear in all the overviews so they are separated out into major and minor characteristics. Table 5 lists these in relation to the overviews in which they appear. Though the nomenclature is different, these results mesh with the themes extracted from the literature on overviews. Below, we address each characteristic in turn.

Attribute/characteristic	**General category**	**Tables of contents**	**File systems**	**Line graphs**	**Abstracts**	
Scope of the material	*descriptive*	x	x	x	x	100%
Quickly guides to information	*exploratory*	x	x	x	x	100%
Shows contextual information	*exploratory*	x	x	x	x	100%
Exposing the structure of the material	*descriptive*	x	x	x	_	75%
Encourages exploration of other information	*exploratory*	x	_	x	x	75%
Provides a snapshot of the state at a particular time	*historical*	x	x	x	_	75%

Table 5: **The attributes and characteristics of an overview based on the analysis in this paper.** The final column shows how often the attribute is represented in the overviews surveyed.

3.1. Showing the scope of the material

Showing scope means that the limits of the detailed data are defined and exposed. This ties in with the theme from the literature that overviews are comprehensive. From the scope, the user knows what they can expect to find in the detail and allows her/him to familiarize her/himself with the whole data set. From this whole, the user can then set filters, as described in [3], to hone in on areas of interest. Only the overview of tabular data [24] provides an overview that is not comprehensive, though this is for usability reasons due to the auditory modality: this gives the user control over the speed of presentation.

3.2. Acting as a guide

Quickly guiding picks up on two themes from the literature: showing salient features and obscuring detail. One reason for overviews is the difficulty of comprehending the whole data set; the abstraction that an overview provides as well as its brevity allows a user to quickly see patterns and relationships that would be harder to see if s/he were perusing the detailed data. By presenting higher level information and doing so in a brief manner, the overview quickly guides to the detailed information. Only from an in depth knowledge of the detailed information could a user glean the sort of understanding that is readily available from an overview.

3.3. Showing contextual information

Having a sense of context allows the user to better understand what is being presented and how it is addressed in the data set. If one wanted to know if some information were present, a search feature would be sufficient. At that point, the detailed data is like a black box where a user dips in to find out some information but does not know what else is there. By showing context, the overview allows the user to familiarize her/himself with the data before dipping in.

3.4. Exposing the structure of the material

The structure or organisation of the detailed information touches on several of the themes from the literature. While context is about what relates to a particular piece of information, structure is about the flow of the detailed information. This can be described as abstraction and obscuring detail in some cases or displaying interrelationships in others (e.g. the structural logic behind the organisation of a file system or a document provides high level information about similarities or disparities between content). However, abstracts do not describe structure as it is not necessary for its purpose and exposing structure would interfere with its prose format. In other words, while structural information is key, it is secondary to the main purpose of an overview which is to best describe the detailed data.

3.5. Encouraging exploration

One of the main purposes of an overview is to act as a guide; encouraging exploration is a corollary to this. The overview should make it possible to discover what is there: not only what the user is looking for, but also what else is there. This ties in with one of the less mentioned attributes of overviews from the literature and also is a minor characteristic per the analysis: in other words, it is desirable but not obligatory. For example, a well-written abstract will entice a reader, while a less well-written one will represent the detailed data but not necessarily engage the readers curiosity.

3.6. Providing historical states

When a data set is dynamic, such as a file system or a working document, an overview can capture the state of the data at a particular time. Tracking changes through the detailed data can be cumbersome and the overall meaning of those changes difficult to understand. Whereas an overview provides an easy way to capture the general, if not the specific changes to the data. This characteristic was not discussed in the literature; this could be because this application of an overview was not of interest in the context of the research. For example, in the overview of personal histories [7], they were interested in patterns in the histories, not displaying their evolution.

3.7. Levels of overviews

So far, we have discussed overviews in general. The aim was to understand overviews as a concept. However, once overviews are used in practice, the context can have a great effect and can muddy the waters. For example, a table of contents describes a text and a portion of the table of contents can describe a subsection of that text. The six characteristics state that an overview should be comprehensive: comprehensive of the text in question. It therefore follows that there can be sub-overviews.

3.7.1. When is an overview a sub-overview?

Logically, a sub-overview provides a comprehensive overview of a subset of a larger dataset. For example, a table of contents could describe a book and a sub-overview might describe a chapter. The sub-overview might take a different form than in the table of contents as the chapter becomes the whole instead of being a subset of the book. In other words, the sub-overview becomes an overview in its own right. Defining it as a sub-overview is only necessary when discussing both the overview of the book and the chapter at the same time.

3.8. Task dependency

The discussion of sub-overviews above highlights that different overviews support different tasks. Looking at a paragraph may only be useful in particular context; conversely, understanding an entire book may be useless in another. In other words, the task and the overview need to be closely related. This task dependancy is included in the six characteristics. This is most obvious in the case of quickly guiding and encouraging exploration: neither of these can be accomplished well without the overview being appropriate for the task in question. Most of the overviews analysed here are general and are fairly modular. They are applicable to a variety of tasks but are not suited to all tasks involving the detailed data they describe. For example, both tables of contents and abstracts can describe a text but they do not serve the same purpose. If the task is to understand the thesis of a text, a table of contents may provide hints but not as well as an abstract could. A successful general overview does a better job of satisfying all the characteristics if the overview anticipates the tasks for which it will be used. A city map could help in several tasks like route-finding, understanding the layout of a neighbourhood or understanding the network of arteries through and around the city. Graphs, on the other hand, are harder to make universal and are often fine-tuned to a task. Another example is a timeline which shows a sequence of events but not interrelationships between them. However, in the case of the timeline, trying to make it more general might compromise its quality: highlighting connections between events or related people could obscure the sequence of events. A timeline quickly guides to the when but not the how. If the task is understanding the former, then the user is more likely to be encouraged to explore. Satisfying the six characteristics aids in assuring that the overview is suitable and useful. Examining the task and overview together and checking, in particular, for the more interactive characteristics, helps determine if the overview is the correct one for the task.

4. AUDITORY OVERVIEWS

We now examine how well our analysis of the characteristics of visual overviews applies to overviews specifically developed for the auditory domain. Table 6 summarises how they match against the auditory overviews in the following analysis.

4.1. U.S. census data

Zhao [25, 26, 8] applies the information-seeking mantra to her work on auditory exploration of U.S. census data. The mantra becomes Auditory Information-Seeking Actions (AISA). Here, in an attempt to distance herself from using visual language, the term *gist* replaces the term overview. A gist is a short audio clip that describes the detailed data. In her thesis, [8] sets forth guidelines for the duration of a gist, its interaction and its latency, giving a more concrete understanding of what a gist or auditory overview should sound like. She argues that a gist should be no more than 10 seconds long due to the capacity of human short term memory, should be low latency (less than 100 milliseconds) and should be synchronised with other modalities to support multi-modal interaction. Only the proposed length of the gist is overview-specific; the other two guidelines pertain to general auditory interaction. Further, gists may contain sub-gists which may be auditioned independently.

The overview of U.S. census data shows four of the characteristics: all of the major ones and one of the minor ones. The overview describes the scope by sonifying data points for all 50 states. The overview is also brief (less than 10 seconds) and acts as a guide by highlighting the variations in population across the United States. Context is heard by listening to neighbouring states. A listener can hear a snapshot of various census by selecting a different census year. It is less clear in the literature whether the overview exposes structure or encourages exploration though it is presumed that large changes in adjacent values might encourage exploration, as might perceived patterns.

4.2. Tabular data

Tabular data is data that is displayed on a grid. This sort of information is difficult to display in audio. In his work sonifying tabular data, [12, 27] do not formally address what an overview is. However, the authors state that in an overview, detail is irrelevant and that an overview can bring out patterns/trends in the data. One key aspect of the tabular data overviews is that they are row/column based. In other words, the overviews cover a subset of the data set. A user gains an understanding of the whole by comparing the row or column subsets. Using a stylus, the user iterates through the columns or rows, controlling the speed of the overview. The stylus interactions also allow the user to focus only on what they feel is relevant as opposed to the entire dataset. As the stylus travels over a column or row, it plays a representation of the numerical data contained within. Thus, [12] do not consider that overviews need to be comprehensive, merely that they represent a large enough subsection that a user can begin to locate salient features. In this case, the data is numerical and the overviews facilitate locating outliers: where numbers are especially high or low.

The tabular data overview shows evidence of all but one – a minor one – of the overview characteristics. The representation of all the columns/rows shows the scope of the detailed data. The brevity of the overview is determined by the user and is also driven by the number of rows/columns. However, the Sonification of each row/column is extremely brief and the overview as a whole can be considered short. The overview guides the users to columns/rows showing high or low values and shows context through contrast to neighbouring rows/columns. The structure is clear as the tabular nature of the data is intrinsic to the overview. The exposition of salient features, in this case high and low values, can encourage exploration. The playing of the overview is user-directed and as a result, it is less clear how well it could represent changes in states (i.e. historical snapshot).

4.3. Edge/node graphs

Edge/node graphs are a way of representing interconnected data. For such graphs, [11] aim to create an auditory equivalent to a

	Scope	Guides quickly	Context	Structure	Encourages exploration	Snapshot
Census data	x	x	x	?	?	x
Tabular data	x	x	x	x	x	?
Edge/node graphs	x	x	x	x	?	x
Source code	x	?	x	x	?	x
Mathematical equations	x	x	x	x	x	x
situate	?	?	x	x	x	x

Table 6: **The attributes and characteristics of overviews from the literature.** The final line of the table shows how many of the overviews reviewed showed evidence of each attribute.

glance. They set out two requirements for the overview: to give an impression of size and complexity, and to describe the topology. The audio glance is an organised iteration through the graph that spreads from the left-most node: each node plays, then each node connected to it and so on. The basis for this is highlighting the relationships between the nodes and not the spatial layout.

The overview of edge/node graphs shows all but one – again a minor one – of the overview characteristics. The overview's main purpose is to describe entire graphs and their layout and thus it exposes the scope and structure of the detailed data. It shows context by describing the interconnections. While the size of the graph drives the length of the overview, it uses short non-speech sounds and is likely to be brief. The exposition of the features of the graph guide the listener through its layout and can provide a snapshot of the graph in various states. There is no evidence to the contrary, but it is hard to determine if the overview encourages exploration.

4.4. Source code

Source code is computer programming code. It is plain text and very syntactically strict. Often, a single code file will have several thousand lines of code. [10]'s overviews describe Java source code. Similar to [11] and [12], they do not address what makes an overview. However, as with Kildal, the focus is on abstraction of the data and detailed data is obscured. The authors concentrate on the types of statements in the code rather than the statements themselves. The code is divided into three categories of statements, with nested statements exposed through a more complex representation. By describing nested statements, a user can perceive the hierarchical structure of the code. In other words, the overview is a broad iteration, describing the entire program and its structure.

This overview is based on [28] who created auralisations of computer programs. They used musical constructs to aid novice programmers to identify bugs in Pascal code. The difference between their work and Finlayson's work is that the goal of [28]'s project, called CAITLIN, was to identify where problems occurred such as improperly terminated IF statements rather than providing a representation of the overall program structure.

Most of the major and minor overview characteristics are present in the overview of program source code. The overview shows scope and structure by summarising the code in order. Context is also apparent through this iteration through the code. Additionally, it is possible to have historical snapshots of the code as it evolves. What is less clear is if the overview is brief. By identifying various programming structures, the overview can act as a guide. However, the length of the code will drive the overview length and code source files can be several thousands of lines long. It is difficult to tell if the overview encourages exploration.

4.5. Mathematical equations

Mathematical equations are the language of maths. They provide complete descriptions of potentially complex relationships between variables and allow for their manipulation and analysis. When read aloud, they can easily and quickly become incomprehensible and misunderstood. [29, 30, 31] tackle the problem of making mathematics more accessible to visually impaired people. They created an auditory glance with the goal of expressing high-level structure to facilitate planning how to approach the mathematical expression. The auditory glance, which is an overview of the equation, describes the general shape of the expression and provides enough specifics to understand the complexity but the specific terms are obscured. For example, a user might hear that something was a number but not what that number is. The auditory glance would allow the user to understand perhaps that the expression is a quadratic equation and the user would need to explore further in order to hear the exact terms of the equation. They used algebraic earcons, composed with timbres, rhythm and prosody to describe the equations. The earcons describe the syntax of the expressionshowing aspects such as super/subscripts and describing the location of the various parts of the equation and their relative sizes. In essence, the auditory glance provides a framework for further exploring the equation. Experiments confirmed that participants could discern the complexity and shape of the equations and they were able to decipher the expressions while listening.

The type of information exposed in the glance shows what [29] prioritised for use in an overview. The specifics were not exposed but the intention was that participants could identify major segments of the expression through prosody and that the glance describes the entire expression. In terms of the characteristics identified through the visual overviews, the auditory glance satisfies all of them. The glance is comprehensive and shows the scope of the expression and describes the location of items through prosody. New items are distinguished through timing and pitch, allowing for context to be heard. The structure and syntax was proven to be discernable through user studies. While not the stated goal of the glance, it could be used to present a historical snapshot. For example, as a user worked with an expression and manipulated its shape, the glance could expose the changing shape by comparing two glances. The stated goal of the glance is to allow listeners to plan how they will approach the mathematical expression and thus it encourages exploration.

4.6. Voice access to web pages

Webpages rely on their spatial layout and visual characteristics to guide the users. Various graphical elements draw the user's eyes

to salient features. However, without a visual component, it is difficult to fluidly navigate the webpage. Here, an overview can facilitate navigation. [9] adapt the Information Seeking Mantra to guide the development of a voice system for accessing web pages. Because of the difference in auditory versus visual interaction, the authors propose *situate* instead of *overview* and describe it as a method to provide "an understanding of [the page's] structure." [p. 857] and to help users locate themselves within the information space. In other words, *situate* answers where the user is on the page and what options are available. The authors imply that this support enables quick navigation from one major section of the page to another. This is not an overview in the strictest sense; it is a system that exposes high level information. However, [9] focus on navigating the detail while occasionally accessing structural information. This is the opposite of some of the other work described here (e.g. [32, 24, 11, 10, 29] etc.) where the interaction starts with the overview which guides users to areas of interest.

The *situate* command in the voice access to web pages [9] is similar to an overview. Its stated goal was to facilitate navigation and expose the structure of webpages. As such, it appears to have several of the overview characteristics. As a navigation tool, it encourages exploration and acts as a brief guide to the page. It is unclear if the overview is semantic or simply structural. This means it is difficult to tell if context is exposed or snapshots of the page are possible. Scope is also difficult to determine.

4.7. How the discovered characteristics fit auditory overviews

Table 6 shows a summary of all the auditory overviews from the literature and how they satisfy the attributes identified through the described survey of overviews. With the visual overviews, there was a clear demarcation between major and minor characteristics. Showing scope, context and quickly guiding all emerged as important characteristics. The remainder (showing structure, encouraging exploration and providing a historical snapshot) were not omnipresent and thus are minor characteristics. This pattern is not repeated with the auditory overviews. Only showing context appeared definitively in all the overviews. Showing scope, structure and providing a historical snapshot were quite frequent while quickly guiding and encouraging exploration were harder to show in the overviews.

The use of the word *definitively* is key here. The overviews surveyed were all easily generalised due to their prevalent natures. Examples of tables of contents, file managers, line graphs and abstracts are plentiful. This is not the case with the auditory overviews where we are reliant on the quality and the comprehensiveness of the written descriptions of original research. This makes it difficult to properly assess them in the same way as visual overviews. As such, this review of auditory overviews is conservative as to which attributes match and which do not.

With this caveat in mind, Table 6 shows that the auditory overviews each match four or more of the six overview attributes. This indicates that the overview attributes do apply to auditory overviews. As far as which attributes apply to auditory overviews, encouraging exploration is the weakest, with only two specifically mentioning navigation and exploration. Encouraging exploration is also the hardest attribute to prove since overview researchers do not address the issue. Guiding quickly is the weakest of the major characteristics, turning up in only four of the six auditory overviews. Overall, the strongest statement that may be made is that the overview attributes seem to apply to auditory overviews as well however, this cannot be proved conclusively.

5. TOWARDS AN AGENDA FOR AUDITORY OVERVIEW RESEARCH

Within the context and themes of the current meeting, it is worth noting that the difficulties inherent in the above analysis would be substantially offset in a research environment which was more conducive to repeatability and transparency. How much easier would the above analysis become if all of the software and data used to implement and test the described auditory overviews were readily available. Many of the points in the above analysis where we are led to infer qualities of an approach or draw weak conclusions concerning the presence or otherwise of a particular characteristic in an auditory overview would become testable, greatly increasing the possibility of making clear comparisons and drawing firmer conclusions.

The identified characteristics however do form a good basis on which to examine the design of any given auditory overview, because they capture the essence of what an overview should aim to do, and this becomes key when the display medium is audio. The fact that sound requires time to audition is both a strength and a weakness of auditory displays. It means that sound is an excellent medium in which to present phenomena that evolve over time, as in, for example, an auditory progress bar, but the obvious drawback is the time required for the display to be heard. It is essential then that auditory overviews avoid wasting time, and enable the user to focus in on any subarea of interest as quickly as possible. Interactivity can have an important role to play here, as for example in the auditory overview of numerical tabular data by [12], where the overview consists of multiple suboverviews which are navigated by the user.

Auditory overviews baring the characteristics identified above have the potential to be pivotal in conveying an understanding of the structure and general characteristics of large data spaces, informing the later exploration of the detail, helping to focus the interaction on points of interest and highlighting salient features and relationships. As the interest in the exploration of large data spaces continues to grow, including "New human-machine interfacing for exploring data (beyond keyboard, mouse and screen)"[4] there appears to be a unique opportunity to make the case for the role of sonification in general, and the use of auditory overviews in particular, in the exploration of large data spaces.

Spatial sound may well have an important role to play here, with its ability to orient the user relative to the different parts of a complex auditory display, and afford the parallel presentation of multiple data sources, as in the work on the Clique system [33] which uses a conversation metaphor – employing up to four spatialised concurrent voices – and task-based interaction to provide auditory access to a Graphical User Interface (GUI) for visually impaired (VI) users. While this work was not done in the context of data exploration, it demonstrates how spatial sound can be used to present multiple concurrent information sources without overloading the user.

The increased power of tablets and mobile phones means that in turn these are being employed for tasks that would previously have required a computer. A limiting factor in their use however remains the small amount of available screen space. Audi-

[4] https://connect.innovateuk.org/web/data-exploration

tory overviews of any kinds of large data spaces, such as maps, spreadsheets, large documents and databases could exploit the fact that audio can be presented over a large virtual space using headphones, again leveraging the properties of overviews described above - providing a complete but rapid presentation of the whole data space, guiding to areas of interest and identifying points of saliency.

There are some situations in which even if the amount of information to be presented is not large, auditory overviews can have an important part to play. Screen-readers currently provide very little by way of overview or summary information to VI users. JAWS[5] includes an "overview" of web pages that is triggered by a keyboard command. This feature lists the number of links, headers, forms and frames on a web page, but in terms of giving the user an idea of the content, this listing of elements provides little more than an indication of how busy the page is. This listing of elements also does not provide any spatial layout – which would aid in collaboration with sighted users (e.g. is there a column layout? Where is the main navigation?) – or any indication of dynamic elements - which remain a source of confusion and are mainly hidden from VI users until they are activated.

Web pages are by no means the only area where auditory overviews have the potential to considerably improve screen reader-based Human-Computer Interaction. Overviews of entire web sites, including indicators of accessibility, of documents and document collections, of large numbers of emails and calendar entries, including indicators of priorities, are among the many areas where overviews could help to overcome the linearity of speech-based displays and improve the efficiency of interaction.

6. CONCLUSIONS

We have examined the nature of overviews, aiming to bring clarity to what is meant by the term and what they are typically intended to do. We have discussed the characteristics of previously reported auditory overviews from the literature, and as far as that literature permits, tried to put these contributions within the overall context. Finally we have identified several areas in which there appears to be considerable scope to extend work on auditory overviews, and discussed why the attributes of overviews we have identified here provide a useful yardstick to guide these future developments.

7. REFERENCES

[1] Richard Chimera and Ben Shneiderman, "An exploratory evaluation of three interfaces for browsing large hierarchical tables of contents," *ACM Trans. Inf. Syst.*, vol. 12, no. 4, pp. 383–406, 1994.

[2] Ben Shneiderman, "Dynamic queries for visual information seeking," *IEEE Software*, vol. 11, pp. 70–77, 1994.

[3] Ben Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," in *Proc. IEEE Conference on Visual Languages*, Boulder, CO, USA, September 1996, pp. 336–343.

[4] Jack Kustanowitz and Ben Shneiderman, "Meaningful presentations of photo libraries: rationale and applications of bi-level radial quantum layouts," in *Proceedings of the ACM/IEEE-CS Joint Conference on Digital Libraries*, New York, NY, USA, 2005, pp. 188–196, ACM.

[5] Bill Kules, Jack Kustanowitz, and Ben Shneiderman, "Categorizing web search results into meaningful and stable categories using fast-feature techniques," in *Proceedings of the ACM/IEEE-CS Joint Conference on Digital Libraries*, New York, NY, USA, 2006, pp. 210–219, ACM.

[6] Marc A. Smith, Ben Shneiderman, Natasa Milic-Fraylingand Eduarda Mendes Rodrigues, Vladimir Barash, Cody Dunne, Tony Capone, Adam Perer, and Eric Gleave, "Analyzing (social media) networks with NodeXL," in *Proceedings of the International Conference on Communities and Technologies*, New York, NY, USA, 2009, pp. 255–264, ACM.

[7] Catherine Plaisant, Brett Milash, Anne Rose, Seth Widoff, and Ben Shneiderman, "Lifelines: visualizing personal histories," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 1996, pp. 221–227, ACM.

[8] Haixia Zhao, *Interactive sonification of abstract data: framework, design space, evaluation, and user tool*, Ph.D. thesis, University of Maryland at College Park, College Park, MD, USA, 2006.

[9] Manuel A. Pérez-Quiñones, Robert G. Capra, and Zhiyan Shao, "The ears have it: a task by information structure taxonomy for voice access to web pages," in *Proc. IFIP Interact*. 2003, IOS Press.

[10] J. Louise Finlayson and Chris Mellish, "The "audioview" – providing a glance at java source code," in *Proceedings of the International Conference on Auditory Display*, Ireland, July 6-9 2005, Department of Computer Science and Information Systems, University of Limerick.

[11] Andy Brown, Robert Stevens, and Steve Pettifer, "Audio representation of graphs: a quick look," in *Proceedings of the International Conference on Auditory Display*, UK, June 2006, Department of Computer Science, Queen Mary University of London.

[12] Johan Kildal and Stephen A. Brewster, "Exploratory strategies and procedures to obtain non-visual overviews using tablevis," in *Intl Conf. Disability, Virtual Reality & Assoc. Tech*, 2006, pp. 285–294.

[13] Chris North, Ben Shneiderman, and Catherine Plaisant, "User controlled overviews of an image library: a case study of the visible human," in *Proceedings of the first ACM international conference on Digital Libraries*, New York, NY, USA, 1996, pp. 74–82, ACM Press.

[14] Andy Cockburn, Amy Karlson, and Benjamin B. Bederson, "A review of overview+detail, zooming, and focus+context interfaces," *ACM Comput. Surv.*, vol. 41, no. 1, pp. 1–31, 2008.

[15] Judith Butcher, *Copy-editing: The Cambridge Handbook for Editors, Authors and Publishers*, chapter 7: Preliminary pages, pp. 177–179, Cambridge University Press, 3rd edition, 1992.

[16] S. Joseph Levine, "Writing and presenting your thesis or dissertation," Retrieved on 8 February 2011, from http://www.learnerassociates.net/dissthes/#26, September 2010.

[5] http://www.freedomscientific.com/products/fs/jaws-product-page.asp

[17] Barbara H. Kwasnik, "A descriptive study of the functional components of browsing," in *Engineering for Human-Computer Interaction*, 1992, pp. 191–203.

[18] Janine Barchas, *Graphic design, print culture, and the eighteenth-century novel*, Cambridge University Press, New York, NY, USA, 2003.

[19] Deborah Barreau and Bonnie A. Nardi, "Finding and reminding: file organization from the desktop," *SIGCHI Bull.*, vol. 27, no. 3, pp. 39–43, 1995.

[20] Eric Freeman and David Gelernter, "Lifestreams: a storage model for personal data," *SIGMOD Rec.*, vol. 25, no. 1, pp. 80–86, 1996.

[21] B. J. Hicks, A. Dong, R.Palmer, and H. C. Mcalpine, "Organizing and managing personal electronic files: A mechanical engineer's perspective," *ACM Trans. Inf. Syst.*, vol. 26, no. 4, pp. 1–40, 2008.

[22] Usability First, "Usability first - usability glossary - focus+context | usability first," Retrieved on 28 January 2011, from http://www.usabilityfirst.com/glossary/focuscontext/, January 2011.

[23] Paul Cairns and Anna L. Cox, *Research Methods for Human-Computer Interaction*, Cambridge University Press, New York, NY, USA, 2008.

[24] Johan Kildal and Stephen A. Brewster, "Providing a size-independent overview of non-visual tables," in *Proceedings of the International Conference on Auditory Display*, June 2006.

[25] Haixia Zhao, Catherine Plaisant, Ben Shneiderman, and Ramani Duraiswami, "Sonification of geo-referenced data for auditory information seeking: Design principle and pilot study," in *Proceedings of the International Conference on Auditory Display*, Sydney, Australia, July 6-9 2004.

[26] Haixia Zhao, Ben Shneiderman, and Catherine Plaisant, "iSonic: Interactive data sonification for blind users," Retrieved on 28 January 2011, from http://www.cs.umd.edu/hcil/audiomap, 2006.

[27] Johan Kildal and Stephen A. Brewster, "Non-visual overviews of complex data sets," in *Extended Abstracts of CHI*. 2006, p. 947, ACM Press.

[28] Paul Vickers and James L. Alty, "Musical program auralization: Empirical studies," *ACM Trans. Appl. Percept.*, vol. 2, no. 4, pp. 477–489, 2005.

[29] R. D. Stevens, S. Brewster, P. C. Wright, and A. D. N. Edwards, "Design and evaluation of an auditory glance at algebra for blind readers," Santa Fe, NM, U.S., 1994, International Community for Auditory Display, pp. 21–30, International Community for Auditory Display.

[30] R. D. Stevens, P. C. Wright, A. D. N. Edwards, and S. A. Brewster, "An audio glance at syntactic structure based on spoken form," in *Proceedings of the 5th international conference on Computers helping people with special needs. Part II*, Munich, Germany, Germany, 1996, ICCHP '96, pp. 627–635, R. Oldenbourg Verlag GmbH.

[31] Robert David Stevens, *Principles for the design of auditory interfaces to present complex information to blind people*, Ph.D. thesis, University of York, UK, 1996.

[32] Catherine Plaisant and Ben Shneiderman, "Organization overviews and role management: inspiration for future desktop environments," in *Conference Companion on Human Factors in Computing Systems*, New York, NY, USA, 1995, pp. 419–420, ACM.

[33] P. Parente, "Clique: A conversant, task-based audio display for GUI applications," in *Doctoral Consortium ASSETS'05*, 2005.