Journal of the Text Encoding Initiative

Issue 10 | December 2016 - July 2019
Selected Papers from the 2015 TEI Conference

“It will discourse most eloquent music”: Sonifying Variants of Hamlet

Iain Emsley and David De Roure

Electronic version

URL: <http://journals.openedition.org/jtei/1535>
DOI: 10.4000/jtei.1535
ISSN: 2162-5603

Publisher

TEI Consortium

Electronic reference

Iain Emsley and David De Roure, « “It will discourse most eloquent music”: Sonifying Variants of Hamlet », *Journal of the Text Encoding Initiative* [Online], Issue 10 | December 2016 - July 2019, Online since 24 January 2017, connection on 03 July 2019. URL : <http://journals.openedition.org/jtei/1535> ; DOI : 10.4000/jtei.1535

For this publication a Creative Commons Attribution 4.0 International license has been granted by the author(s) who retain full copyright.

"It will discourse most eloquent music": Sonifying Variants of Hamlet

Iain Emsley and David De Roure

- 1 Sonification is a complementary technique to visualization that uses sound to describe relationships in data. Kramer defines sonification as "the use of nonspeech audio to convey information." More specifically, "sonification is the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation" (Kramer et al. 1999). While providing new opportunities for communicating through the human perceptual and cognitive apparatus, sonification poses challenges with presenting the exploratory patterns in data to the user, as sonification techniques are less well established than those of visualization.
- 2 We describe our work to sonify textual variants¹ in order to aid exploratory textual analysis. The sonification presented focuses on using a mixture of instruments and pitches to help the user listen to differences in the structure between variations of a text or texts encoded in Text Encoding Initiative (TEI) XML.

- 3 Our approach is inspired by the Hinman Collator, an opto-mechanical device originally used to highlight print variants in Shakespeare texts, whereby visual differences between two texts literally stood out through a stereoscopic effect (Hinman 1947, Smith 2000). Using an audio stream for each text, this project aims to produce a binaural presentation of the text, creating an audio version of the stereoscopic illusion used in collating machines. The timing and frequencies can be extracted for storage and transformation into alternate formats or to repeat the analysis.
- 4 We present initial work on XML variants of Shakespeare's *Hamlet* using the XML content from the Bodleian Libraries' *First Folio* project² and their earlier project, the *Shakespeare Quartos Archive*.³ We extracted document entities such as acts, scenes, lines, and stage directions for the analysis. These are viewed as hyperstructures: structures that can be manipulated by hyperoperations to create new structures. These may be separated from the text for sonification and comparison with other variants. Analytical perceptions can be altered through the presentation of the instruments, pitches, and audio icons. Audio displays demand the creator to rethink how structural data is presented to the user, and about the hyperstructures extracted to give potential for conversion of the analysis into hypermedia using visualization as well as sonification. Early results show promise for the auditory comparison.
- 5 We look at related work and present the case study. We then consider the use of audio beacons to help the user locate within the document, and discuss the integration with visualization. Finally, we look at future work and conclude the paper.

1. Related Work

- 6 Sonification on data patterns has been explored in several projects. For example, work on stock market data (Nesbitt and Barrass 2002, Nesbitt and Barrass 2004) discusses the use of volume and pitch to alert to changes in the data, rather than relying on purely visual stimuli. It demonstrates the use of sonification for pattern analysis in exploratory data using a rule system, and is closely associated with visualization.
- 7 The *Listening to Wikipedia* project⁴ presents an audio-visual display of edits made to Wikipedia pages. Using circles for the visualization and rule-based sounds, it presents the "recent changes" feed to the user, including new users and the type of user making the edit. This work provides an elegant interface to the user data but it is limited to one stream.

- 8 The TEI-Comparator⁵ was developed to compare paragraphs and visualize the changes (Cumming and Mittelbach 2010, Lehmann et al. 2010) for the *Holinshed Chronicles*⁶ project, illustrating a collation approach applied to TEI. This visualization work does not render the text into audio signals, and it was designed for a particular text. It focuses on the text rather than the editorial structures.
- 9 Sonification of hyperstructures is explored where an authored hypertextual structure is sonified using the techniques of algorithmic composition (De Roure et al. 2002). In contrast, we present work that develops the notion of sonifying the hyperstructure, or hyperstructures, extracted and transformed from existing editorial matter.

2. Sonifying Versions of *Hamlet*

- 10 We present work on creating an auditory display using Shakespeare's *Hamlet*. This began with the Bodleian's work on the First Folio and continued with their earlier work on the variants of *Hamlet* in the Shakespeare Quartos Archive project with the British Library and Folger Shakespeare Library.
- 11 This work focuses on an alternative presentation to Hinman's Collator. In the collator, two texts are transposed in stereoscope to show the differences between them. Our eyes use variations between images to interpret depth in 3D vision; similarly, our ears use subtle timing and phase variations to establish a stereo stage. Using an audio stream for each text, the project aims to produce a binaural image of the text with auditory beacons to guide the user within the audio illusion. Playing a synchronized audio stream per text in each ear helps the listener's brain to hear any subtle differences between two versions through use of binaural transmission.
- 12 Displaying the hyperstructures of the texts, such as the speakers of a <line> element, allows the listener to hear whether editorial changes have been made to the textual structure. This method uses chosen structures from the metadata to examine the texts and how they might change over time or between editors.
- 13 We convert a selection of TEI XML elements and attributes, including @act and @scene attributes of the <div> tag, <stage>,⁷ lines (both <l> and <p>), and <person>,⁸ into a series of numbers using a rule set to cluster the elements into related groups. Table 1 illustrates the sonification pipeline. The process uses the @xml:id attribute for the characters in the <line> elements to match the <speaker> to the <line>.

- 14 Initially, we locate the speakers to build a representation of marked-up characters and identities. In the P5 TEI XML used in the *First Folio*, we can take this from the headers using the <speaker> element. This provides an id that is used with the variants that may be found in the text. We create a linked list to associate the @xml:id attribute with an incremental numeric value.
- 15 In the quartos, we have to extract the names where the @ref was taken from the <name> element that contains the @character attribute. As the identities are not marked up in the <head> element, we use a linked list to capture the ids marked up in the @character attribute and give these a numeric value as they are discovered. This creates a linked list of identities with an assigned number for that character in that version of the text.

Table 1. Table of the initial rule-based mapping.

Element	XML to MIDI Number	MIDI Number to Sound
act	0–9	Flute
scene	10–20	Flute
stage	40–50	Shakers
speaker (l, p)	60–100 (added to person number)	Flute
Person (speaker or name)	Number derived from association with the @xml:id	

- 16 A simple rule-based mapping was applied to the transformations, as described in [table 1](#), to turn the element into a number using the same groups. This use of rules provides a method of ensuring that the different types of TEI encoding from the Shakespeare Quartos Archive and the First Folio could be mapped to the same groups using numbers in the MIDI (Musical Instrument Digital Interface) specification.

- 17 The generated numbers that represent the desired parts of the editorial markup are then either stored as a file or streamed to the sonification software. This transformation is completed as a separate script from the sonification toolkit. We preserve the mapped data for provenance purposes.
- 18 Our sonification software, written using the `Chuck9` language, ingests the file or the data stream. The data is then mapped to relevant pitches and sounds, using a set of rules that are created from a discussion between the developer and the listeners, as shown in [table 1](#). Some initial sounds were created as experiments, but the choice should be an exercise in codesign. The data is transformed into the frequency to be played with the given instrument, and Chuck defines some instruments that we use. We might use more than one channel if using more than one file, provided the playing device supports this and it is defined in the mapping. The sound is then played. The frequencies created and their associated times are optionally written to a file so that we can check that the data is being presented correctly.

Figure 1. Sample transform of TEI XML structure into sound.

- 19 We show this pipeline in the sequence diagram in [figure 1](#). The transformation script may be written in any language. We use both PHP and Python, having developed the scripts from other software. We use our own software to map the sounds to the notes and then play them.
- 20 The different versions of TEI encodings pose challenges to ensure that each play has the same characters encoded in the same way. Either the markup used changes between the texts or the character order is different and may produce alternate pitches across variations for the same person. This echoes an issue that we face using digitally marked-up copies of texts for sonification:

what are we sonifying—the text or the editorial structure? In this work, we concentrate on looking at the structures that relate to a particular variant, such as a stage direction or person speaking on a particular line, rather than the markup pertaining to the text's physical condition or features.

- 21 By way of example, in the 1603 Quarto edition¹⁰ the first stage direction and first lines are:

```
<stage rend="italic, centred" type="entrance">Enter two Centinels.
  <add place="margin-right" type="note" hand="#ab" resp="#bli">
    <figure>
      <figDesc>Brace.</figDesc>
    </figure>now call'd
    <name type="character" ref="#bar">Bernardo</name> <lb/>&amp;
    <name type="character" ref="#fra">Francisco</name> –
  </add>
</stage>
<sp who="#sen">
  <speaker>1.</speaker>
  <l><c rend="droppedCapital">S</c>Tand: who is that?</l>
</sp>
<sp who="#bar">
  <speaker>2.</speaker>
  <l>Tis I.</l>
</sp>
```

- 22 In the 1605 Quarto edition,¹¹ the stage direction and first lines are:

```
<stage rend="italic, centred" type="entrance">Enter
  <name type="character" ref="#bar">Barnardo</name>, and
  <name type="character" ref="#fra">Francisco</name>, two Centinels.
</stage>
<sp who="#bar">
  <speaker rend="italic">Bar.</speaker>
  <l><c rend="droppedCapital">VV</c>Hose there?</l>
</sp>
<sp who="#fra">
  <speaker rend="italic">Fran.</speaker>
  <l>Nay answere me. Stand and vnfolde your selfe.</l>
</sp>
```

- 23 Although the sentinels are identified as Barnardo and Francisco in the stage direction, the text and markup specify different characters between the variants. In our software, this would create separate sounds for the first line but not the second. The latter line would create the binaural illusion through the production of the same note and volume in both ears so that it appears to be one entity. The former breaks this by changing the note and volume so that it is clearly two sounds representing a variation in the markup. At this point, the listener understands that the editorial choices are dissimilar: the line is marked up differently in the versions being compared.
- 24 This allows us to understand the variant editorial structures placed onto the texts, reflecting choices made either in the encoding or in the textual version. As the structures are rendered at the same time, we hear the differences simultaneously.
- 25 We group related elements, such as <act>s and <scene>s, together through sounds to aid the comprehension of the structure. We are aware that the brain does have a limit to the amount of audio information it can decode, and we try to use the sounds in a way that limits cognitive overload where possible. This grouping is another transformation that creates an artificial layer of interpretation of the editorial choices, designed to help the listener understand the data more easily.

3. Audio Beacons

- 26 The use of sound, with or without visual assistance, poses user experience challenges. The sounds and the relationships that they describe may be unfamiliar and require some training or assistance. As audio is an unfamiliar medium for work in exploratory analysis, there is a need to help the listener identify their position within the document structure and for designers to create an experience that justifies the work.
- 27 As user experiments to represent the sounds we considered a mixture of auditory icons (Gaver 1997), which are sounds that mimic the real world, and "earcons" (Blattner, Sumikawa, and Greenberg 1989), which use music and might be thought of as leitmotifs to indicate a presence or event. The present work uses earcons created from a variety of computer-generated instruments, such as flutes or types of shakers, to represent different types of textual event as sound.

Considering pitch as "a psychological phenomenon related to the frequency" (Levitin 2006), gestalt psychology is used to represent similar events with sounds to help the listener's perception of the interpretation (Rosli and Cabrera 2014).

- 28 The work's acts and scenes provide useful beacons for the listener to understand which section of the text is being presented. Simple auditory icons are used to aid the listener in understanding the presented event, and research is ongoing to improve these. In the present sonifications, we use flutes for both acts and scenes, using lower pitches for the acts and slightly higher ones for the scenes. The intention is twofold: firstly, we use the scenes and acts to mark locations, like ticks on a graph axis, denoting a position; secondly, we group the non-speaking parts of the text that are contained in the data set into similar pitches to help the listener identify the events. Using grouping lessens the cognitive load of identifying the event being presented and echoes the existing practices within visualization.
- 29 In early versions of the sonification, the acts and scenes were produced with different instruments and pitches to allow the user to identify them as part of this group. Currently we represent the two elements with one instrument, the flute, as they are so closely linked as structures, but with a more rapidly rising pitch for the <scene> element.
- 30 The <stage> element provides greater detail to use within the display. The current sounds use shakers to denote the stage element. Different pitches are used to denote the @type attribute so we can hear a subtle difference of type of stage direction. The @type and @who attributes help to design the type of sound. The sounds associated with the @who attribute can be linked to the speakers but present a different issue. The @speaker attribute is associated with one person, but the stage directions may have more than one person interacting with the direction. This changes the sound from a single note to a chord or progression. We use different pitches for the @type attribute so that the change in the type is detectable.
- 31 The @xml:id attribute in the <l> tag is used to link to the speaker where the @xml:id is checked against the list of person associations to retrieve the numeric identity, as per table 1. The flute is used for the speakers and the tone is then matched to the id in the linked list derived from the person list. This has the number 60 added to it so that the MIDI pitch starts near middle C and can be heard easily. The volume for each speaker is slightly raised as they continue speaking, helping the user identify that the speaker has not changed. When comparing two streams, the listeners

will identify any textual changes when both the tone associated with the speaker and volume alter. Using the two parameters of note and volume provides the user with two axes to understand where data changes.

- 32 The `actsspeak` MP3 file demonstrates the acts, scenes, and speakers. The first flute sounds represent the act and then the scene. A standard time-step of 100 milliseconds is applied to other elements, which are silent in this file, and then the shakers sound to indicate the scenes. In the `speakers` MP3 file, the flutes create one note as the texts match (from the Hamlet XML shown above). When the texts diverge, so do the sounds.
- 33 This means that the listener requires a key to understand how to associate the sound with the event. The present version of the software uses simple pitches and instruments. We are considering the development of auditory icons to help identify the type of element event being presented.
- 34 What sound should be represented: one that is contemporary to the text or to the period when the text is being sonified? Even if we use a period sound for representation, we must be aware that we are still creating an artificial sound.
- 35 Using period sounds raises questions of interpretation and performance as well. This interpretive element could be useful to demonstrate the way that sounds might work on a stage for students or to explore the soundscape presented in the play from any existing sounds. This would rely upon knowing various contexts, such as performance and staging practices as well as the locations of performance.
- 36 Early results from informal listening tests suggest that the pitches need to be listenable for the user as part of the user experience. This raises questions about two contexts: the listener and the underlying data. The sounds that are created should be distinguishable from the background noises in their local environment, otherwise the sonification will not be understood. Equally, the sound that is created should work with the underlying data set and the facets being enhanced or created. This echoes the question raised from musicology about the use of historical techniques and instruments in playing pieces and whether these represent the period or an interpretation of it (Holden 2012). These are sound design issues and require reflection on the aims of the sonification or the exploration, as well as using known musical techniques, such as masking or filtering, or the

use of psychoacoustics, allowing the mind to make some of the cognitive links. The sounds used for the elements should be similar, such as using two wind instruments, as two dissimilar sounds, such as a wind and a string instrument, are distracting and uncomfortable for this purpose.

- 37 This allows us to focus upon the use of the sonifications to show the relevant facets. The current objective is to explore differences between the two versions of a particular text, leading us to focus on the moments where the sounds diverge. In other experiments, we use sound to demonstrate the richness or sparsity of data while searching catalogues of metadata, as an adjunct process to a search. As well as showing the results and linking to the data, we provide a view of the way that the data is organized both in volume of data and in its closeness to earlier results.
- 38 At present, our software only allows us to listen to the audio from start to end as one piece. We have not yet added interaction to allow control of the sounds or annotation. This would lead us into a deeper consideration of methodologies and practices.

4. Visualization

- 39 Multimodal experiences provide alternative presentations for the same event. In [figure 1](#), we show an early prototype visualization showing symbolic representations of the events, using the [Processing language](#) used for coding in the visual arts.¹² This visualization was used to investigate how a multimodal experience might enhance the sonification work.
- 40 The note data were sent to a visualization process to show an abstract image or text based on the notes received, displayed in near-real time to the sound. Such images or texts were found to aid comprehension of the audio display.
- 41 The use of abstract symbols, like the circles for speakers, poses the same challenge as in sonification: the symbol must be understood. This confirms that we should use a key to the symbols for easier comprehension.
- 42 Informal tests were achieved by playing the sonifications to interested parties, including a workshop, using our equipment. Notes were made on the feedback for discussion and used as some of the user tests for the software. This suggests that further refinement is required to help make the displays more useful and points towards different use cases. This may include using text and developing a version for the web. Multimodal experiences provide a different challenge to the

user. The pure sonification can be a passive task and run simultaneously to other tasks with the anomalies in the sound change alerting the user to variances. The use of visual cues makes the task active as the symbol must be seen to be effective.

5. Conclusion

- 43 The use of sound opens up possibilities for visually impaired scholars to interpret and explore data sets using sound. A mixture of audio user experience for presentation and haptics for control and annotation via a simple set of tools could provide another way of collating or exploring texts and their markup. As previously discussed, mixing the soundscape of the texts with performance context may provide a different experience of the texts and how they could be perceived.
- 44 The work presented here explores one facet of sound as "the transformation of data relations into perceived relations in an acoustic signal" (Kramer et al. 1999). As a starting point, it shows us how sound can be used as a secondary perceptive detail. The early results are promising and provoke questions about how we can use sonification to illuminate multiple facets of the data.
- 45 We have demonstrated the potential of sonification as a tool to help the user identify differences between textual variants. Although known in exploring scientific data, sonification is a new analytical approach for the Digital Humanities. It allows the designer to use multiple parameters simultaneously to add meaning to an event by changing types of sound, tone, pitch, or volume. Making the technique understandable presents an ongoing challenge. The use of binaural playback indicates that further work with spatial displays to create a richer user display may improve user comprehension of the data.
- 46 Words and lines may be auralized using the tone associated with the speaker. The sonification would then render the associated tones. This does pose the issue of how a word is sonified: is it by length or some other metric? The <choice> element from the Text Encoding Initiative provides options for an original reading and a variation. The sonification would then have to associate a similar tone with the choices. It may be that the original text would be the expected tone and that the variation is an additional different pitch played simultaneously.
- 47 Further work is needed to create better auditory icons that work across streams and to integrate audio and visual displays. We have not explored this area fully. Contextual questions include the type of sound that would be typical in a dramatic context or physical one, such as the construction

of places of performance. It also demands knowledge of the practices of staging or presentation. We intend to conduct formal user testing. This has implications for the development of sonic skills for both listeners and developers to provide a good user experience.

- 48 While these initial examples are focused upon individual works, sound may also be used within discovery processes. We are also looking at social network analysis within given timeframes to explore how a community interacts, and this may raise questions regarding the linking of different types of event and text genres, such as a novel or letter, or whether it has a license and is available or not. We have applied sound to the searching of authors within the [EEBO/TCP metadata catalogue](#)¹³ to explore the richness or sparsity of data within a search with an option to show links to the resulting texts.
- 49 We believe that the use of sound provides an exciting workflow for exploring hyperstructures to determine differences between them. The novelty in this area is a major challenge, but we strongly believe that it will be useful in the exploration of variants between texts marked up with TEI.

BIBLIOGRAPHY

- Blattner, Meera M., Denise A. Sumikawa, and Robert M. Greenberg. 1989. "Earcons and Icons: Their Structure and Common Design Principles." *Human-Computer Interaction* 4 (1): 11–44. doi:10.1207/s15327051hci0401_1.
- Cummings, James, and Arno Mittelbach. 2010. "The Holinshed Project: Comparing and Linking Two Editions of Holinshed's Chronicle." *International Journal of Humanities and Arts Computing* 4 (1–2): 39–53. doi:10.3366/ijhac.2011.0006.
- De Roure, David C., Don G. Cruickshank, Danus T. Michaelides, Kevin R. Page, and Mark J. Weal. 2002. "On Hyperstructure and Musical Structure." In *Proceedings of the Thirteenth ACM Conference on Hypertext and Hypermedia*, 95–104. NY: ACM. doi:10.1145/513338.513366.
- Gaver, William W. 1997. "Auditory Interfaces." In *Handbook of Human-computer Interaction*, 2nd ed., edited by Martin G. Helander, Thomas K. Landauer, and Prasad V. Prabh, 1003–1041. Amsterdam: Elsevier Science, 1997.
- Hinman, Charlton. 1947. "Mechanized Collation: A Preliminary Report." *Papers of the Bibliographical Society of America* 41 (2): 99–106.

- Holden, Claire. 2012. "Recreating Early 19th-century Style in a 21st-century Marketplace: An Orchestral Violinist's Perspective." Paper presented at Institute of Musical Research DeNote Seminar, Senate House, London, January 30: 17–21 http://orca.cf.ac.uk/17241/1/Claire_Holden_IMR_Seminar_doc.pdf.
- Kramer, Gregory. 1993. *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Reading, MA: Perseus Publishing.
- Kramer, Gregory, Bruce Walker, Terri Bonebright, Perry Cook, John H. Flowers, Nadine Miner, John Neuhoff, et al. 1999. "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. Santa Fe, NM: International Community for Auditory Display (ICAD). <http://www.icad.org/websiteV2.0/References/nsf.html>.
- Lehmann, Lasse, Arno Mittelbach, James Cummings, Christoph Rensing, and Ralf Steinmetz. 2010. "Automatic Detection and Visualisation of Overlap for Tracking of Information Flow." In *Proceedings of I-KNOW 10: 10th International Conference on Knowledge Management and Knowledge Technologies*, edited by Klaus Tochtermann and Hermann Maurer, 186–97. Graz, Austria: Verlag der Technischen Universität Graz. <http://hdl.handle.net/10419/44446>.
- Levitin, Daniel. 2006. *This is Your Brain on Music: Understanding a Human Obsession*. London: Atlantic Books.
- Nesbitt, Keith V., and Stephen Barrass. 2002. "Evaluation of a Multimodal Sonification and Visualisation of Depth of Market Stock Data." In *Proceedings of the 2002 International Conference on Auditory Display (ICAD 2002)*, edited by Ryohei Nakatsu and Hideki Kawahara, 2–5. <http://hdl.handle.net/1853/51355>.
- . 2004. "Finding Trading Patterns in Stock Market Data." *IEEE Computer Graphics and Applications* 24 (5): 45–55. doi:10.1109/MCG.2004.28.
- Rosli, Muhammad Hafiz Wan, and Andrés Cabrera. 2014. "Application of Gestalt Principles to Multimodal Data Representation." In *Proceedings of the IEEE VIS Arts Program (VISAP) 2014*, 102–107.
- Smith, Steven E. 2000. "'The Eternal Verities Verified': Charlton Hinman and the Roots of Mechanical Collation." *Studies in Bibliography* 53: 129–62.
- TEI Consortium. 2013. *TEI P5: Guidelines for Electronic Text Encoding and Interchange*. Version 2.5.0. Last updated July 26. N.p.: TEI Consortium. <http://www.tei-c.org/Vault/P5/2.5.0/doc/tei-p5-doc/en/html/>.

NOTES

- 1** Iain Emsley, Scripts for sonification of the TEI XML, 2015, accessed December 22, 2016, <https://github.com/iaine/sonificationshakespeare>.

- 2 *The Bodleian First Folio: A Digital Facsimile of the First Folio of Shakespeare's Plays*, Bodleian Arch. G c.7, accessed December 22, 2016, <http://firstfolio.bodleian.ox.ac.uk/>.
- 3 The Shakespeare Quartos Archive, accessed December 22, 2016, <http://quartos.org/>.
- 4 Accessed December 22, 2016, <http://listen.hatnote.com/>.
- 5 Accessed December 22, 2016, <http://tei-comparator.sourceforge.net/>.
- 6 "About the Project," accessed December 22, 2016, <http://www.cems.ox.ac.uk/holinshed/about.shtml>.
- 7 TEI Consortium 2013.
- 8 TEI Consortium 2013.
- 9 Chuck: Strongly-timed, Concurrent, and On-the-fly Music Programming Language, accessed December 22, 2016, <http://chuck.cs.princeton.edu/>.
- 10 *The Tragedy of Hamlet Prince of Denmarke: An Electronic Edition*, First Quarto, 1603, British Library Shelfmark: C.34.k.1, accessed December 22, 2016, http://www.quartos.org/XML_Orig/ham-1603-22275x-bli-c01_orig.xml.
- 11 *The Tragedy of Hamlet Prince of Denmarke: An Electronic Edition*, Second Quarto Variant, 1605, British Library Shelfmark: C.34.k.2, accessed December 22, 2016, http://www.quartos.org/XML_Orig/ham-1605-22276a-bli-c01_orig.xml.
- 12 The Processing Language, <https://processing.org/>, accessed December 22, 2016.
- 13 Catalogue of the Text Creation Partnership / Early English Books Online, <https://github.com/textcreationpartnership/Texts>.

ABSTRACT

Sonification is a complementary technique to visualization that uses sound to describe relationships in data. We describe work to aid exploratory textual analysis by sonifying textual variants. The sonification presented focuses on using pitch and tones to help the user listen to differences in the structure between variations of a text or texts encoded in Text Encoding Initiative (TEI) XML. Extracting hyperstructures, we describe our conversion of TEI elements and attributes into sounds for a listener. We discuss our approaches to creating the sounds used to represent the data from the Bodleian Libraries' *First Folio* project and early visualizations, and we consider the issues raised by the use of this novel technique. The use of sound provides an exciting

alternative way of exploring textual structures to determine differences between them. While the novelty in this area is a major challenge, we suggest that this method can be useful in the exploration of variants between texts marked up with TEI.

INDEX

Keywords: auditory display, hypertext, sonification

AUTHORS

IAIN EMSLEY

Iain Emsley is a research associate at the Oxford e-Research Centre. Currently reading for a Master's in Software Engineering at Oxford, his research interests include sonification.

DAVID DE ROURE

David De Roure is Professor of e-Research at the University of Oxford, where he directs the multidisciplinary e-Research Centre. Focused on advancing digital scholarship, David has conducted research across disciplines in the areas of social machines, computational musicology, web science, social computing, and hypertext.