THE MSPACE CLASSICAL MUSIC EXPLORER

Improving Access to the Classical Music Space for Real People

mc schraefel**, Daniel Alexander Smith,
Alistair Russel, Alistair Owens, Craig Harris, Max Wilson
info @ mspace.fm http://mspace.fm

Overview

The mSpace Classical Music Explorer, available at http://demo.mspace.fm, is a Web-based explorer designed to improve access to the classical music domain, especially for those who do not have domain knowledge. We could characterize the approach as designed to support those who might say "I don't know much about classical music, but I know what I like when I hear it." For this group, browsing a list of domain terms is not helpful, because those terms do not have meaning for the group; key word search is equally unhelpful, since again the group does not yet know what they wish to specify. The mSpace classical music explorer addresses the access problem by providing a variety of direct manipulation techniques that enables exploration and selection without domain expertise. In the following paper, we describe these techniques. We also present results of evaluation informing the multimedia aspects of the interface, and briefly describe the mSpace model on which the explorer is based.

Figure 1. mSpace, showing preview cue open in the Composer dimension.

mSpace Explorer Features: Improving Access

The mSpace Explorer supports a variety of interactive techniques to facilitate access to the domain. These include the direct manipulations supported by the interface, the mechanism of associating information in the layout, and the use of multimedia to enhance perception of the domain. Videos demonstrating these operations are available at the mSpace website, http://mspace.fm.

The mSpace Explorer consists of a multicolumn layout (Figure 1). Each column represents a dimension in the domain, like era, composer, genre, arrangement/instrument, and so on. A selection in a left column filters what will then be displayed in the column/dimension to the left. Any time an instance in a column is selected (like Romantic in Era), information describing that instance is displayed in the info pane below the columns, and the column to the right is populated with information appropriate to that constraint (like composers in the Romantic period). Associated with each instance in a column is a *preview cue* set. By highlighting the speaker icon next to an instance (Figure 3), a set of several pieces associated with that instance is made available; the first cue in the list begins to play; the others can be selected at any time, and can be "scrubbed" to any location in the piece. At any point in the interface if the person wishes to keep track of any element they've discovered in their exploration, they can double click the element to have it added to a Favorites area. The preview cue remains associated with this favorite.

While the explorer presents a default set of dimensions, a person can add, subtract or rearrange dimensions: if it is more meaningful to see the classical music world organized by instrument first rather than composer, and then by recording artist and producer rather than genre, or followed by genre, the interface facilitates these choices (Figure 6). We now describe the design rationale and evaluations informing the current approach to the mSpace explorer. We follow this with a brief overview of the software framework used in the application.

Preview Cues: Will that be one Cue or Two?

One of the critical elements of the mSpace explorer is the preview cue [schraefel 03, schraefel 04], in this case, an audio preview cue. The audio preview cue lets a person rapidly sample parts of the domain to see which areas are of interest for further exploration. Unlike many online music store services which only support 20 second samples of a single piece, preview cues in the mSpace Explorer approach are provides as a lightweight, rapid mechanism to support *exploration* of the domain. The association of a cue with an instance in a dimension means that a person with no specific knowledge of what that selected term means can still perform what Marshal describe as "infomation triage" [Marshall 94]: where none was possible before, at least a binary judgement can now be made about whether or not the current selection is of interest to them to explore further. Such triaging is a powerful tool: a means of access is provided to a previously inaccessible space. In previous work [schraefel 03], we describe how we tested preview cues in a variety of layout conditions (Figures 2 and 5). We found that the association of a single preview cue with each instance in a domain enhanced search effectiveness, efficiency and satisfaction. More particularly, to use Dillon's evaluation approach of *process, outcome, affect* [Dillon 01] where affect is the perceived experience of using the software, participants reported feeling empowered in their searches when preview cues were present (it was difficult to pull some people away from the trial).

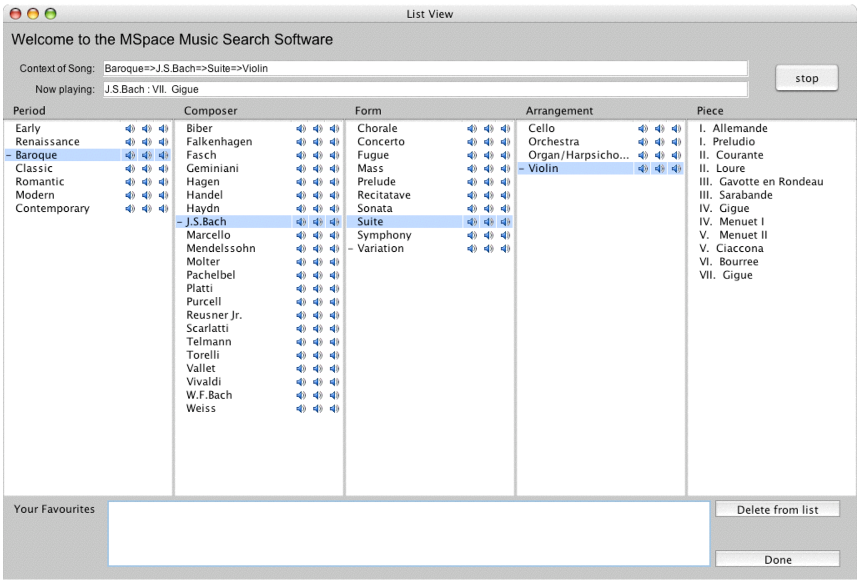

Figure 2. Testing preview cues: a spatial layout with multiple cues per instance

A question that came out of this work was whether or not one cue was a sufficiently representative sample of a given part of a domain. In follow up studies [schraefel 04] we found that there was no difference in performance whether one cue or three were used. What we found, however, was that search patterns changed: participants tended to go deeper into a selected area with multiple cues, whereas they did more rapid triaging of a wider area of a domain with a single cue available. Multiple cues were also used less the deeper into the domain the participant went.

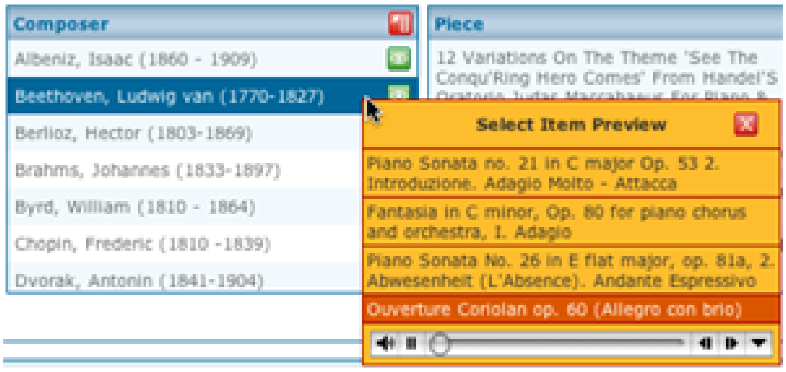

Figure 3. mSpace version 0.4 preview cues

Preview Cue Selection. There are multiple models for preview cue selection: random selection from the set of properly constrained pieces (the selection Romantic | Beethoven would not include preview cues of Hyden); matches with expert lists for the selected instance; community based recommendation (other listeners who have chosen Romantic have selected these pieces for their favorites); or, when the technology is more advanced, feature extraction (computationally extract the representative features from the set of pieces that matches a specific criteria, and select the closest matches to these).

We have found no evidence that any technique is better than any other for helping people access the classical music domain. On the contrary, it seems that *ANY* audio clue to an area is simply, but powerfully and effectively, better than none. Associating preview cues with instances is a relatively low cost process that yields high access returns.

Layout

As a complement to preview cues, the mSpace explorer uses a *spatial* multicolumn layout (Figures 1 and 5) rather than the *temporal* layout (Figure 4) still common on the web. In a spatial layout the results of selections made in one part of the display are displayed in another part of the same display, *not* on a different page that occludes the previous page. This persistent display assists a person to maintain awareness of surrounding or *contextual* information, while reducing the need to remember what went before the current selection.   Clicking through Yahoo's category searches is an example of a temporal display: clicking on one category opens a new page, erasing the previous page from view. Only the path to that page is persistently available. In developing the current multicolumn approach, we tested both the Yahoo-like temporal approach and the multicolumn spatial approach and found that performance was significantly better with the spatial layout [schraefel 03]. Indeed, we found a negative correlation between age and layout: whereas performance remained strong for all age ranges with the spatial approach, there was a noticeable drop off as age increased (50+) with the temporal browser.

Figure 4. Testing a single column (temporal) version of the mSpace Classical Music Explorer: path information is available at the top of the column.

Slices

One of the motivations for the mSpace approach is to let people explore a domain space from a perspective that suits their needs and interests. To support this flexible kind of interaction, the mSpace explorer assumes that a space like classical music has many dimensions. We are not very good at imagining multidimensional spaces. Indeed, 25% of the population cannot manage 3D displays represented on 2D screens [Modjeska 00]. This finding was also part of our motivation for keeping our browser as flat as possible. Indeed current research has shown that 3D is most effective for animation transitions between 2D states , rather than replacing 2D representations completely [Robertson 02].

Figure 5.A spatial layout: leverages recognition rather than recall. In this case, hovering a cursor over an instance initiated preview cue playback

Thus a challenge of the mSpace explorer is to find a way to represent richly dimensional spaces in less than 3D. Our approach has been to use the analogy of a plane: if a transparent plane were inserted into a multidimensional site to be used as a view perspective, the available dimensions could appear flattened onto that surface, creating the appearance of (temporary) hierarchies. mSpace replicates the plane effect by creating *slices* in a dimensional space. The explorer presents a default starting slice. In the case of the classical music explorer the default slice is three dimensions: Era | Composer | Piece. Several operations on slices are available: sorting, adding and subtracting.

Sorting. Slices can be rearranged by moving columns to the left or the right of each other.
Adding enables additional dimensions to be added to a slice.
Subtracting, conversely, enables dimensions to be removed from a slice.

These interaction operations mean that the multidimensional space can be sliced dynamically in whatever way best suits the needs or interests of the person exploring the space. Such manipulations also act to foreground relationships within a space: it becomes possible to explore associations for instance between era and the popularity of a given form or arrangement, or between style and country of composition or of composer.

Favorites: Supporting Information Triage

To enhance support of information triage, we have the Favorites area: an instance of interest can be added to the Favorites area. The preview cue set associated with the instance stays with the instance when added to a favorite. Two operations are possible in favorites: (1) by hovering over the preview cue icon, people can rapidly remind themselves of why they may have selected the cue; (2) by clicking on the instance, the path taken (dimensions selected) to record that instance is represented in the interface. Facilitating information triage with a quick selection area means that a person can focus on exploring the domain, rather than having to remember what they wanted to look at further. The favorites area provides specific task support for user-determined focus: exploration of new resources and relationships is unimpeded, while recovery and further investigation of discovered resources is facilitated.

The mSpace model and Software Framework

The mSpace Classical Music Explorer application is built on the open source mSpace software framework, version 0.4 (http://mspace.sourceforge.org), described in detail in [Harris 04]. The framework likewise provides the interaction attributes described above - slices, sorting, adding, subtracting, info views and preview cues - so that this interface approach can be wrapped around any data source. Thus the generic mSpace framework can thus be wrapped around any data to create a domain oriented mSpace explorer. We have also deployed the interface component of the framework using javascript to provide the interface manipulations. We use javascript because it is a browser-interpreted language, and therefore does not require a person to download any additional code for the application to be useful.

While it is beyond the scope of this paper to describe the architecture of the framework, suffice it to say that it is enabled by Semantic Web [Berners-Lee 01] technologies. The semantic mechanisms within the system allow for automated processes to discover and associate new information to enhance the domain. So for instance, the current Classical Music application does not have a dimension on country of composer, or on historical events of periods. The semantics means, however, that given a discovery process, the Explorer will be able to find and automatically associate this new dimension so that people could add into a slice the new History dimension next to Composer, for instance, to see historical events that occurred within the Composer's life time.

Figure 6. mSpace Classical Music Explorer, showing controls to add new dimensions beside an existing dimension.

Augmenting Information Views: There is no single Beethoven

The current version of the mSpace Classical Music Explorer associates information about instances in a dimension in a rather manual way: sources on classical music from the Web are identified by a researcher or expert. We use a set of scripts to process that information into Semantic Web form (converting the data to RDF [RDF 04] and storing it in a triple store (Rusher 01, Harris 03)), and from there the information - images, text, video - is associated automatically with instances as they are selected.

The current approach to information views gives the impression that there is a single way to describe an instance: a selected composer yields an image of the composer and some descriptive text. There is more than one view, however, on a composer - one school may see a composer as a genius; another as a mean spirited crank of little talent; there is more than one context - historical, musical, biographical - in which to consider a composer. Hypertext research has long celebrated the possibility of linking readers with multiple versions of information in a domain via something called a "distributed link base" [Herzog 95] where annotated links - links with notes describing the content they represent = from different locations are automatically harvested from distributed sources to be made available for exploration of an associated topic. Within hypertext research, however, distributed link services required considerable human intervention, and usually worked in a small scale. The Semantic Web makes distributed link bases possible to generate both automatically, and as Web scale.

Data Sources

In the case of the mSpace Classical Music Explorer, we created seed test data set by starting with a large collection of MP3s. Their ID3 tags [ID3] contained metadata about the recording; these tags were augmented manually with information about the composer, the dates of the composer's life, the form of the work and the instrumentation about the arrangement. This data was then converted into RDF. The metadata categories were then available to be used as dimensions in the mSpace. The information sources to associate with the instances in the categories were discovered manually from the Web, and then *scraped* to be processed likewise into RDF so that the appropriate data could be automatically associated with the instances on demand as selected.

The mSpace Interaction Design Model

The mSpace framework is based on the mSpace interaction model, described in [Gibbins 04, McGuffin 04]. The mSpace model is a formal description of the properties and behaviours of an mSpace. Formally, an mSpace is a special case of a graph structure known as a polyarchy [McGuffin 04; Figure 7). This description is not particularly interesting in the specific context of the Classical Music Explorer application. By having a formal description of how relationships between parts of an mspace work, however, we make it possible to compare the mSpace model against other kinds of graphs structures for modeling information. This allows us to explore comparatively the affordances and limitations provided by a given approach. Such comparisons enable software designers to make choices about which graph structure may be better than another for presenting the kinds of relationships in the data they wish to support.

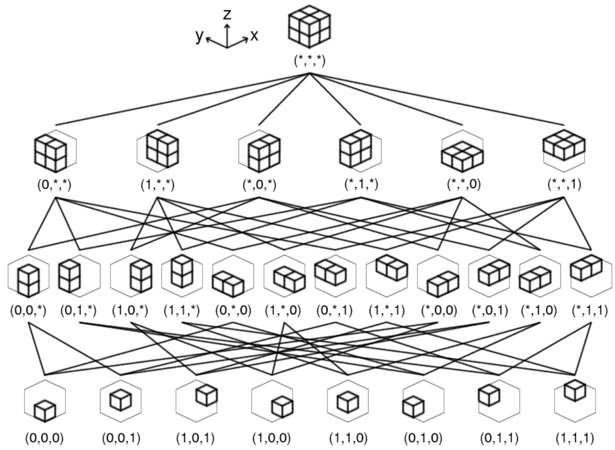

*Figure 7. A small 2*3 mSpace polyarchy, showing the possible slices and selections.*

Likewise, the mSpace model acts as a structure for specifying how to embody mSpaces in any kind of implementation. We have chosen to develop the mSpace model within a Semantic Web framework to make it easy to deploy an mSpace against Semantic Web protocols. We have done this for several reasons: (1) to leverage the Web scale of data that can be available to associate with a domain; (2) to enable other features made possible by the Semantic Web, in particular, inference. The Semantic Web's use of ontologies [Fensel 00] to describe both information entities and the relationships between them (for example, symphony has composer) enables a Semantic Web application to be able to generate data that is not explicitly expressed in the data itself. For instance, it would be possible to extrapolate a set of composers who are early adopters of a new genre or style of music by looking for earliest examples of a form and their associated composers. It would likewise also be possible to infer the popularity of an artist over time by looking at concert programs published in newspapers, for instance, to see how frequently or not certain works have been performed, or similarly, to identify specific points of popularity or decline. (3) The Semantic Web makes it simple to *type* data. By knowing data type (number, string, date, map coordinates) we can provide multiple representations for the information. The mSpace framework takes advantage of data typing: it makes it possible to represent dimensions as charts or on maps, so that for instance it would be possible to have a map populated with concentrations of composers, over time, so that by selecting periods one could see if one region of the world more so than another was hot at a given time. Color might be used in such a map to indicate concurrently who had the most output, or whether works were religious or secular, or which works where were

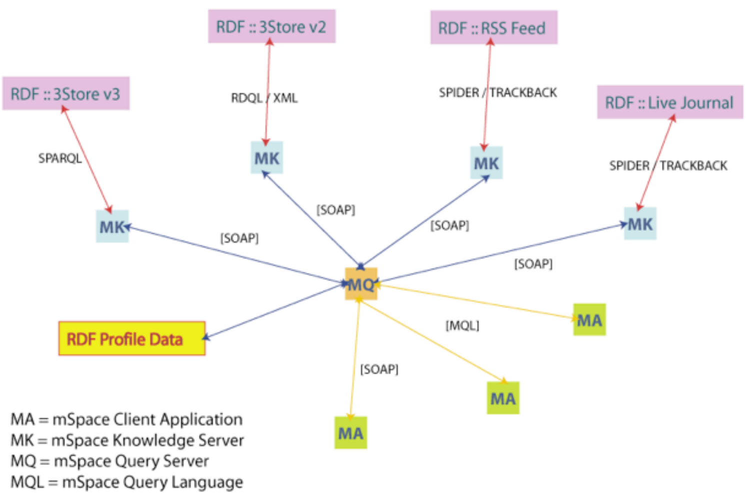

Figure 8. Revised mSpace architecture to support automatic discovery of resources

patronized by whom. A picture is worth a thousand words, a dynamic picture that can show relationships or change, perhaps a thousand or so more.

Future Work

We are currently developing a publishing and data discovery architecture (Figure 8) to use Semantic Web protocols to automate the discovery and capture of relevant new information to be associated with Information views - effectively generating link bases for domain instances [schraefel 05]. Part of this process is the development of a discovery/publishing process which enables people creating information to point their information at an mSpace (in this case, classical music) Query server and have that information become instantly available to be associated with the instance in a dimension. For instance, if someone publishes in a blog or web page or other source of information about Beethoven's relationship to Napoleon, that source can become part of the possible contexts available for a person to explore.

We are also developing a service to allow different aggregated mSpaces to integrate common attributes, and allow inter-mSpace browsing. One mSpace will be able to connect to another by this common attribute, and will offer the user a greater selection of additional attributes at that point. For example, when browsing a movie mSpace (such as the Internet Movie Database (IMdB)), one may look at the soundtracks of a particular film. This may offer the composers of these songs, and a person may wish to know more about the composer, and, by capture of the associated Classical Music mSpace, be able to learn more about the composer's particular works, which classical era they belong to, and with associated preview cues, be able to hear more of that composer's work.

Conclusion

In this paper, we have presented an overview of the mSpace Classical Music Explorer application. We have described the interaction features of the application and the rationale and supporting research for the interface design. We have also presented an overview of the mSpace software framework, the open source general framework to enable developers to create mSpaces for any information domain. While we have described the manual process of data discovery we used to populate the classical music application, we have also pointed to ongoing research to enable automatic discovery and publishing of data sources to augment the current application.

The main focus of our work has been to develop a model and supporting framework to make it easier for application developers to improve access to information domains where people have an interest but not the expertise to make use of current common Web search tools like key word search. As presented here, our research has shown that the combination of effective layout and associated, usable information like preview cues in particular, enhances people's experience and sense of ability to access a domain. In the case of classical music, our success has been to provide an effective tool to enable people to access, explore and discover new knowledge and understanding about a domain of interest that had previously been experienced as inaccessible.

Acknowledgments

The mSpace project is partially sponsored by the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Consortium, EPSRC, UK (http://www.aktors.org), and by the School of Electronics and Computer Science, University of Southampton (http://www.ecs.soton.ac.uk). The first evaluative studies of Preview Cues were supported by the Center for Innovative Technology Ontario (CITO).

References

[ID3]. Id3v2 made easy. id3.org: Last Accessed: March 2005. http://www.id3.org/easy.html.

[Berners-Lee 01] Berners-Lee, T., Hendler, J. and Lasilia, O. The semantic web. Scientific American, 2001.

[Dillon 01] Dillon, A. Beyond Usability. Canadian Journal of Information Science 26(4) (2001), 57-69

[Fensel 00] Fensel, D. Ontologies: Their glory and the new bottlenecks they create. OntoWeb Workshop.Semantic Web Project Proposal., Vrije Universiteit Amsterdam (the Netherlands), 2000.

[Gibbins 04] Gibbins, N., Harris, S., Dix, A. and schraefel, m.c. Applying mSpace interfaces to the semantic web - working paper. IAM Group, University of Southampton: 2004.

[Harris C 04] Harris, C., Owens, A., Russel, A. and Smith, D.A. Mspace: Exploring the semantic web. A technical report in support of the mspace software framework. IAM Group, University of Southampton: 2004.

[Harris S 03] Harris, S.a.G., N., 3store: Efficient bulk rdf storage. In Proc 1st International Workshop on Practical and Scalable Semantic Systems, (Sanibel Island, Florida., 2003), 1-15.

[Herzog, 95] Herzog, M. (1995) "The Dexter Hypertext Reference Model. http://www.dbai.tuwien.ac.at/staff/herzog/thesis/node25.html

[Mashall 94] Marshall, C.C.S., F.M., Coombs, J.H., Viki: Spatial hypertext supporting emergent structure. In Proc ACM European Conf. on Hypermedia Technology, (1994), ACM Press, 13–23.

[McGuffin 04] McGuffin, M.J. and schraefel, m.c., A comparison of hyperstructures: Zzstructures, mspaces, and polyarchies. In Proc Fifteenth ACM conference on Hypertext & hypermedia, (Santa Cruz, CA, USA, 2004), ACM Press, 153-162.

[Modjeska 00] Modjeska, D. Hierarchical data visualization in desktop virtual reality. PhD Dissertation, Department of Computer Science, Univeristy of Toronto. Toronto: 2000.

[RDF 04] RDF Primer, W3C Technical Reports and Publications, 2004. http://www.w3.org/TR/rdf-primer/

[Robertson 02] G. Robertson, K. Cameron, M. Czerwinski, and D. Robbins. Polyarchy visualization: visualizing multiple intersecting hierarchies. In Proceedings of ACM CHI 2002 Conference on Human Factors in Computing Systems, pages 423{430, 2002.

[Rusher 01] Rusher, J. Triple store. (2001). Last Accessed: http://www.w3.org/2001/sw/Europe/events/20031113-storage/positions/rusher.html.

[**schraefel 03] schraefel, m.c., Karam, M. and Zhao, S., Listen to the music: Audio preview cues for exploration of online music. In Proc Interact 2003 - Bringing the Bits Together, (Zurich, Switzerland., 2003).
A longer version of this paper, with more study detail is available as schraefel, m. c., Karam, M. and Zhao, S. (2003) Audio Preview Cues: Support for Exploration of Music Information Spaces. Technical Report ECSTR-IAM03-004, Electronics and Computer Science, University of Southampton (2003).

[schraefel 04] schraefel, m. c., Wilson, M. and Karam, M. Preview Cues: Enhancing Access to Multimedia Content. Technical Report, ECS, University of Southampton, (2004).

[schraefel 05] schraefel, m. c., Smith, D. A., Owens, A., Russell, A. and Harris, C. (2005) The evolving mSpace platform: leveraging the Semantic Web on the Trail of the Memex. ACM Hypertext, 2005, Salzburg, *in press*.

*** The use of lowercase in "schraefel" is deliberate (http://web.uvic.ca/wguide/Pages/Capitals.html)*