

BROWN UNIVERSITY

Garibaldi & LADS

Interactive Multitouch Systems for the
Visual Arts

Donnie Kendall

5/18/2011

Table of Contents

- Background 2
 - Motivation..... 2
 - Garibaldi on the Surface 2
 - Reception 3
- LADS 3
 - Design Overview 3
 - Back End Design 4
 - Modules 6
 - Artwork Selection..... 6
 - Artwork Mode..... 7
 - Knowledge Web 8
- Funded By 9

Background

Motivation

In 2005, the Garibaldi Panorama was donated to the Brown University Library by Dr. James Smith. This unique work, a double-sided watercolor panorama measuring 300 feet in length and four and a half feet tall, presented many display challenges, and was placed in storage until an effective solution was found.

The panorama was digitized and made available as a series of web pages in 2007. Unfortunately, this still precluded the artifact from exhibition in either library or museum because of its unusual dimensions, dual-sided presentation and delicacy.

In 2009, Professor Andries van Dam was approached by the Brown Library. They had heard of the Microsoft Surface and wanted to see if something could be developed that would allow library visitors to interact with the panorama in a tactile fashion within the library, in lieu of the internet. From this meeting, the *Garibaldi on the Surface* project was born.

Garibaldi on the Surface

Garibaldi on the Surface was designed to bring the Garibaldi Panorama out of Brown Library vault, allowing museum patrons and other visitors a chance to view and interact with this unusual piece. The Microsoft Surface was targeted because of its large form factor combined with multitouch input, a unique at the time. We wanted the experience of virtual interaction with the panorama to be as close as possible to working with the tangible object.

In addition to the standard gestures of pinch-zooming and swiping left and right already standard on the Surface, we allowed for a few distinctive interactions; chief among these was the virtual magnifying glass. This option afforded a user the ability to place a virtual magnifying glass on the Surface to zoom



Figure 2 - Garibaldi on the Surface UI



Figure 1 - Garibaldi on the Surface with its external display

in on a small portion of the virtual Panorama without using the entire screen necessitated with the pinch-zoom gesture. This interaction presented users with a magnified view of one part of the Panorama to examine small details while maintaining the overall visual context lost during a full-screen pinch-zoom.

We also provided the option to create *hotspots* on the Panorama, providing additional context-sensitive information about certain sections. This allowed for items, such as correspondence and images from newspapers contemporary with the Panorama, to be embedded within the piece in an overall context. These documents could be shared on a large external connected to the Surface via a flick gesture. The ability to embed videos and to play back scene-specific narrations in multiple languages was also provided.

Reception

Garibaldi on the Surface was widely well-received as a successful blending of computer science with the humanities. Microsoft decided to continue funding the project in order for us to expand on the current functionality. The British Library featured it in their exhibition “Growing Knowledge: The Evolution of Research” which ran from October 2010 through July 2011, and showcased new technologies supporting academic and research goals.

LADS

The expanded functionality came about as LADS - Large Artwork Displayed on the Surface. Our overall design has moved beyond being Surface-specific; we target primarily large-format touch devices but have the ability for the software to run on most Windows 7 touch-enabled devices that support at least two-contact multitouch. In widening this scope our hope is that LADS will be used by museums outside of Brown that may not own a Surface nor have the budget to acquire one, as well as allowing it to run on very large touch devices.



Figure 3 - LADS on a 60" touch display

Design Overview

Garibaldi on the Surface was designed around the display of only one artwork: the Garibaldi Panorama. The primary goal of LADS was to move beyond this, allowing users to explore any of the provided

catalogue of artwork. These pieces are generally expected to include large format scrolls, panoramas, or tapestries; however, any two-dimensional artwork with a high-resolution scan could be viewed.

Internal changes were made to the data file to accommodate this increase in breadth, and an additional artwork selection mode was created to allow users to select the artwork that they wanted to explore.

In artwork mode hotspots were retained. Several simple image manipulation tools were added, allowing users to change options, such as the brightness or contrast, of the artwork to help emphasize hidden details. In addition, there are now documents associated with the artwork as a whole, instead of just with individual hotspots. Users have the option to save these documents in a personal workspace to browse later.

An additional *knowledge web* was created, allowing users to examine the interconnections between artwork. These connections can be created by keyword tags on each artwork; from here users could explore a semantic web of interrelated artwork.

Back End Design

The back end consists of a series of XML files and a large data directory containing the images. These images are stored at several different resolutions, ranging from small thumbnails to Deep Zoom image pyramids, to enable viewing at a variety of sizes and detail levels.

For the XML files there is one master XML file which contains references to all of the artwork along with the keywords associated with each artwork and other relevant details. This is the first XML file which LADS reads upon being launched, and from it other relevant data is loaded.

```

<Image path="path.bmp" title="Title" artist="artist" medium="medium" year="1127" longitude="116°23'E" latitude="39°55'N">
  <Keywords>
    <Keyword Value="Bianjing" />
    <Keyword Value="Kaifeng" />
    <Keyword Value="Neo-Confucianism" />
    <Keyword Value="Qingming Festival" />
    <Keyword Value="shan shui style" />
    <Keyword Value="Song period" />
    <Keyword Value="Taoism" />
  </Keywords>
  <Locations>
    <Purchase longitude="120.2" latitude="30.2" />
    <Display>
      <Location longitude="121.5" latitude="25" />
      <Location longitude="116.45" latitude="40" />
    </Display>
    <Work longitude="113.5" latitude="34.75" />
  </Locations>
  <Metadata>
    <Group name="A">
      <Item Filename="item.jpg" />
      <Item Filename="litem2.jpg" />
    </Group>
  </Metadata>
</Image>

```

Figure 4 - Example XML entry for a single artwork

An example of the artwork XML fragment is given in Figure 4. This block could be repeated as many times as necessary for the various artwork present in the catalog. The *image* section contains primary details – the filename of the artwork, its title, artist name, and the current location of it in latitude and longitude. The *keywords* block contains the keywords associated with the artwork for the knowledge web. The *Locations* block contains locations other than its current display location which are of relevance – locations at which it was previously displayed, sold, painted at, etc. The *metadata* block could be repeated up to 4 times (A, B, C, D) and allowed for metadata to be display in knowledge web mode.

In addition to the overall XML, each artwork has an XML documented associated with it to populate its hotspots. An example of this XML file is shown in Figure 5. The *hotspot* section is repeated as many times as necessary to populate the hotspots for an artwork. For each hotspot the *title* and location in X and Y coordinates is specified. The *type* can be either text or image, and *description* either contains the filename of the image or the text block to be displayed.

```

<hotspots>
  <hotspot>
    <name> Sky </name>
    <positionX> 2000 </positionX>
    <positionY> 4000 </positionY>
    <type> text </type>
    <description> description </description>
  </hotspot>
</hotspots>

```

Figure 5 - Artwork XML Example

Modules

Artwork Selection



Figure 6 - Artwork Mode

Artwork selection mode is the gateway into the LADS system. Shown in Figure 6, it has a great deal of functionality packed into a simple UI. Along the middle of the screen is a timeline, with relevant historical events displayed under the years which they span. Below that the artwork in the catalog is displayed, which each artwork centered under the year associated with it in its *Image* entry in the main XML file. When multiple pieces are closely associated in years they are stacked and tiled so that each is as close as possible to the appropriate year. This view can be zoomed with a pinch-zoom gesture, enabling the artwork to spread out in order to more correctly match their year of completion.

When an artwork is selected it is highlighted with a yellow frame, and a larger thumbnail is displayed to the right of the world map. On the world map the *Locations* from the XML file are displayed as color coded dots. Tapping on any of these indicates the action that occurred at that location. To the left of the world map more detailed information is displayed. Tapping on the thumbnail preview to the right of the map selects that artwork and takes the user to the main artwork mode.

Tapping on the filter button below the map brings up a dialog allowing the user to filter artwork based on Artist, Location, or Keyword. This allows users looking for a specific artwork to rapidly locate it even in busy catalogs. Selecting a region on the map further filters artwork to only those relevant to that region, allowing users to concentrate on a region of interest if they so choose.

Artwork Mode



Figure 7 - Artwork Mode

The main goal of artwork mode is to provide a relatively uncluttered interface for exploring the artwork itself. As seen in Figure 7, the largest part of the space is occupied by the artwork itself. This can be expanded by tapping the bars to the right of and above the left sidebar and bottom bar, causing them to collapse to the screen edge and providing maximum space for viewing of the artwork.

The artwork itself is displayed from a Deep Zoom image pyramid, allowing artwork scanned at a



Figure 8 - Artwork mode hotspots

sufficiently high resolution to be zoomed-in close enough to see individual brush strokes, and zoomed-out far enough to see the entire artwork on large panoramas.

The left sidebar contains the various controls for interacting with the artwork. At the top are three sliders for temporarily

adjusting the brightness, contrast, and saturation of the artwork. This allows users to get a better view of near-hidden details in the art. Below that is the hotspots interface. Tapping on the *hotspots on* button displays the hotspots as a series of location specific dots. Tapping on a dot displays the information associated with it (see Figure 7).

Below the hotspots interface is the associated media interface. Tapping on an entry in this displays a large version of the media in question allowing users to view contemporary newspapers, hand written notes, and other artifacts of interest. These can be dragged to the bottom bar and saved for later viewing.

In the bottom left hand corner is a small thumbnail representation of the artwork being viewed. This is covered by a highlighted area indicating the current view area of the main artwork display. As the user zooms in and out this is updated allowing them to see how large of an area of the artwork they are currently viewing.

Knowledge Web

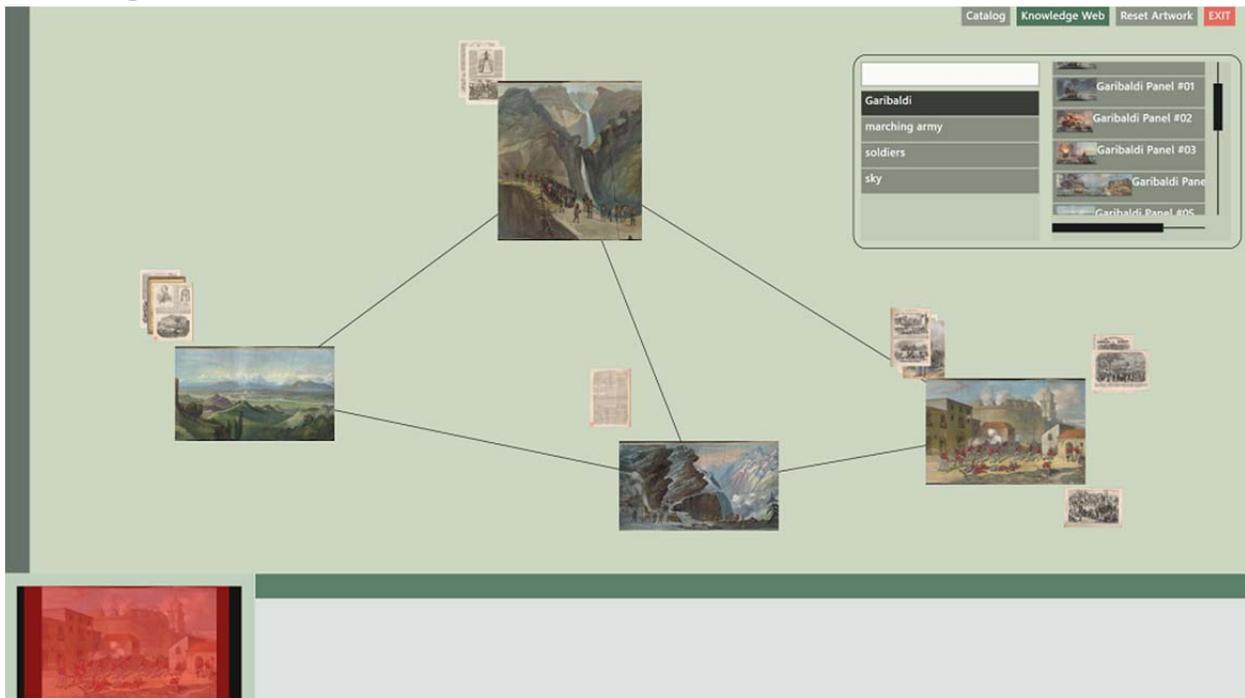


Figure 9 - Knowledge web mode

Knowledge web mode (shown in Figure 9) allows users to explore the interrelations between the artwork in the catalog. When a user first enters this mode the artwork that they were viewing in artwork mode is centered in the screen. By touching and holding on this artwork a search dialog appears showing the keywords associated with the artwork. By typing they can search through this list, with the results filtered in real-time. When a keyword is selected a list of all other artwork sharing that keyword is shown along with thumbnails.

The user can add another artwork to the web by selecting it from this list. An added artwork has relation lines drawn from it to all of the other artwork in the web which share that keyword. From here

users may then hold on the new artwork to get a list of its specific keywords and continue building the relationship web.

At any point the user can return to artwork mode by tapping on any of the artwork present in the web. This takes them to artwork mode with that particular artwork loaded, allowing them to explore it in full detail and see all of its associated documents.

The documents associated with each artwork are shown in piles around the artwork. These piles are controlled by the XML file, allowing the system to present like documents grouped together. For example – if there are a large number of newspaper articles as well as some hand drawn sketches, these could be shown in two different piles. Tapping on a pile expands it so that each document in it can be seen. Tapping on one of these documents then selects that document, bringing it to full size and allowing full interaction with it.

Funded By

Funding for these projects was provided in part by Microsoft Research, the British Library, and the Brown Library.