Brown University

LADS and TAG

Touch-first Systems for Museum Exhibition and Display

Alexander Hills
5/9/2012

Table of Contents

Background

In 2005, the Garibaldi Panorama was donated to the Brown University Library by Dr. James Smith. The piece is 300 foot long, four and a half feet tall panorama. Its size and unusual aspect ratio presented many challenges, and it was placed in storage.

In 2007, the panorama was digitized and turned into a large series of web pages. While this did make the panorama available for a broad audience, it didn't manage to connote the scale or true dimensions of the piece. Additionally, there was no interface designed for library or museum exhibitions themselves; the created application was only suitable for web viewing.

Two years later, Professor Andries van Dam and the Brown Library began work together to create an application utilizing the Microsoft Surface that was suitable for displaying the entire panorama in its historic, artistic, and cultural contexts. The application would use the touch-focused interaction promoted by the Surface to display the panorama in an intuitive, intimate fashion. From these plans, the Garibaldi on the Surface project was created.

Figure 1 – Garibaldi on the Surface UI

Figure 2 – Garibaldi on the Surface in test with its secondary display

Upon completion of the project, the potential of large touchscreens to effectively display enormous and unique artworks became apparent. Garibaldi on the Surface was designed for a single piece of artwork - it worked well for one piece, and displayed it effectively. However, while there are a number of museum exhibitions that are focused around a single piece of art, far more are centered on a collection of pieces, some subset of the museum's catalog.

LADS – Evolution of Garibaldi on the Surface

The future of the Garibaldi project was established in LADS (Large Artwork Displayed on the

Surface). The primary goal of LADS was to move beyond the previous projects' limitation of one artwork and view an entire collection of artworks in full detail, along with their relevant metadata and context. Additionally, where the original project was designed for the Microsoft Surface and ran exclusively on the Surface, LADS runs on any Windows 7 computer from a high-end workstation to a low-end desktop. LADS is optimized for large screen touch interfaces, but all manipulations are possible with just a mouse.

User Experience Walkthrough

Figure 3 – Catalog Mode

Catalog Mode

LADS opens to the catalog mode, shown in Figure 3, an innovative viewer for the collection contained within the LADS database. The catalog mode contains a few separate elements: the timeline, the map, and the detail areas. The timeline is a zoomable, panable arrangement of all the imagery in the collection by time, dynamically arranged to display artworks as close to their respective times of creation as possible while maintaining their relative positions in space. This allows for an at-a-glance view of the layout of artworks over the established timeframe. Zooming and panning are both accomplished via established touch gestures (pinch-zoom and swiping).

Artworks can be filtered by any criteria by expanding the filter box (done by tapping the bar labeled "Filter"). This brings up a selection area that allows one to dynamically change the content visible in the timeline by simply selecting specific criteria. This refreshes the layout on the bottom, changing positions of artworks and the scale of the timeline to accommodate the new content.

Additional information about any individual artwork can be displayed by tapping an artwork to select it. Textual metadata is displayed in the detail area on the top left; this data can be any set of key-value pairs, but by default are Title, Artist, Year, Medium, and Keywords (which is a list of values). Additionally, upon selection of a piece, fine-grained geographic information is displayed on the map in the center. Hotspots appear on the map in the form of color-coded circles; each hotspot can either show the location of creation, display, or current home of the artwork. A supplemental large thumbnail is displayed in the top-right as well, allowing for easy recollection of what image is selected, as well as allowing the viewer to see the image in more detail. Tapping this image sends the user to artwork mode.

Artwork Mode
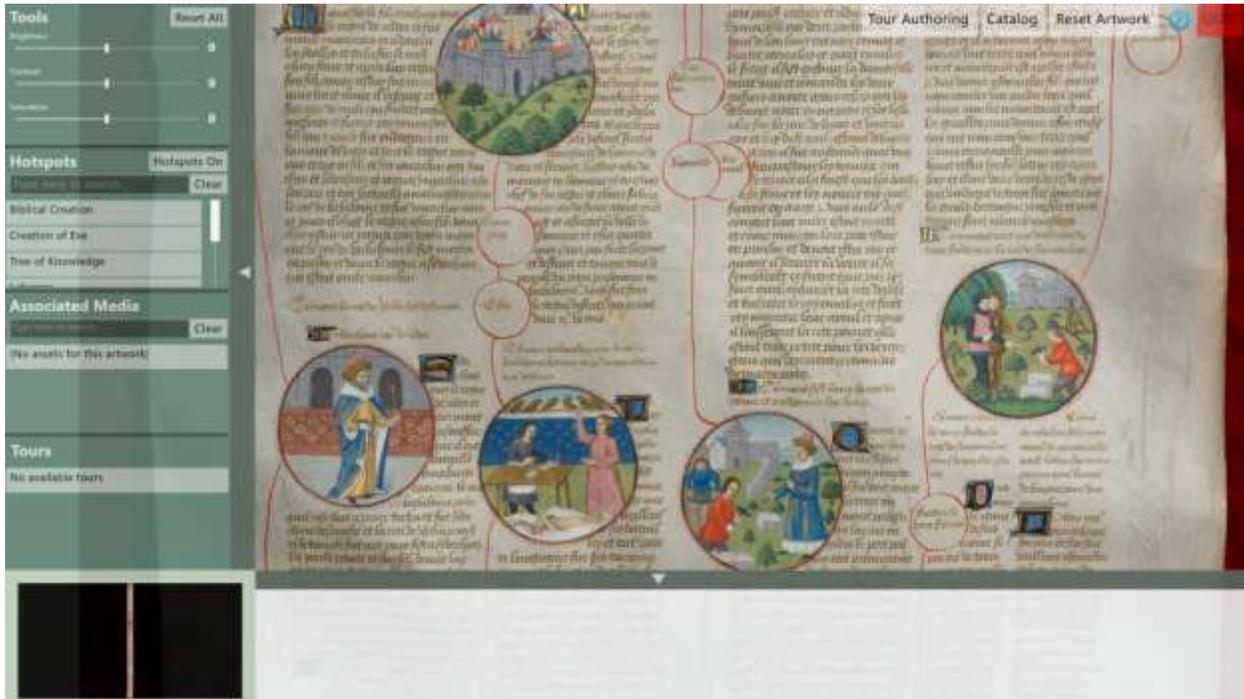
Figure 4 – Artwork Mode

Artwork mode is the core of LADS – it's where large artworks are viewed in full detail at maximum resolution. As can be seen in Figure 4, the vast majority of the screen is taken up by the artwork itself, with toolbars to the left and bottom of the UI. Both toolbars can be docked to the side and bottom by tapping on the rectangle or arrow at the inward edge of the bar itself, providing the largest possible viewing area for the piece.
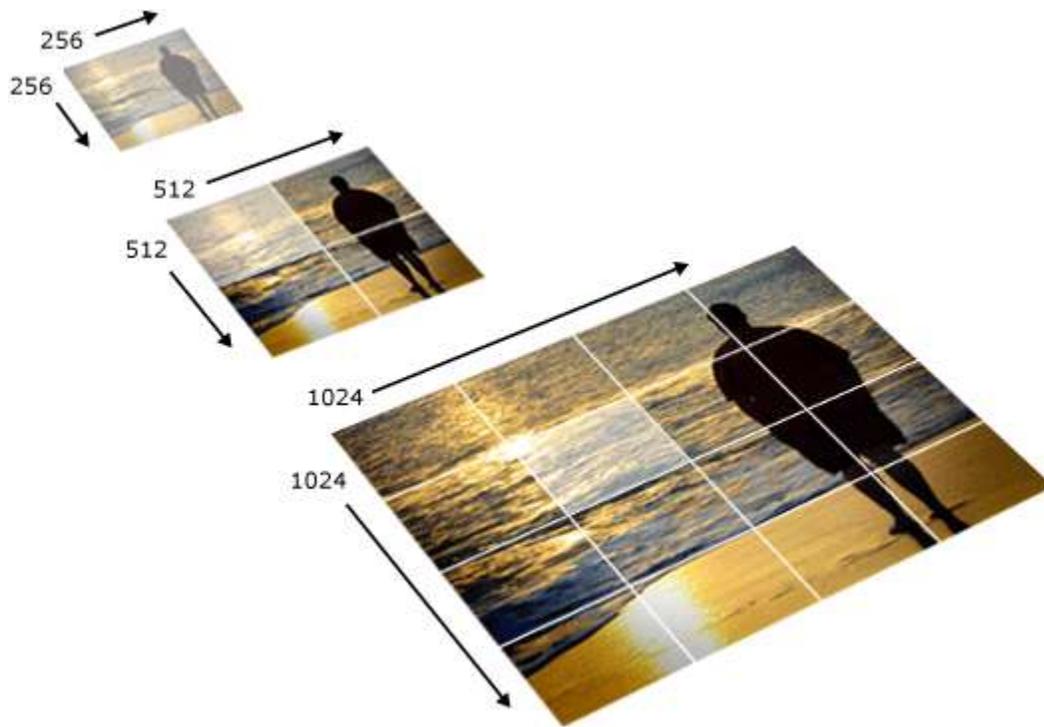
Figure 5 – DeepZoom pyramid

LADS uses an open source implementation of Microsoft's DeepZoom technology specifically made for Windows Presentation Foundation, instead of its native Silverlight. DeepZoom enables an application to display any sized image at full resolution, no matter what scale. It accomplishes this task by creating an image pyramid, beginning at 1x1 pixels and extending to the dimensions of the piece, then tiling each level of the image pyramid and saving the tiles to the disk (seen in Figure 5). This allows only the necessary portions of an image to be loaded at once, allowing any section of a gigapixel image to be viewed at full resolution in real-time without having to load the entire image into ram (this is necessary for enormous images, for instance the Garibaldi Panorama which clocks in at 10 billion pixels or approximately 37 gigabytes).

This technology is used in Artwork Mode to display the artwork itself, allowing artworks of arbitrary size to be examined. This is an incredibly useful tool, and enables a viewer to see detail down to the brushstroke level of a piece, as can be seen in Figure 6. The ability to examine brushstrokes and the process an artist used to create a piece allows a viewer to examine many artitsic qualities of a piece that could not be otherwise seen in a standard museum context.

Figure 6 – Brushstrokes in the Garibaldi Panorama

There are a number of other tools available in artwork mode that allow the user to view the

Figure 7 – Associated Media

historical and cultural context of the piece. The first of these is called associated media, as demonstrated to the left. Associated media are images or videos that are related to the current artwork in a general way, such as the newspaper clipping that the piece is based off of or the initial sketch the artist made as a study. These assets are WPF Surface SDK Scatterview Items, and as such inherit many touch-focused behaviors. They can all be pinch-zoomed, dragged, and tossed around, complete with inertia/momentum for a realistic user experience. Assets are easily disposed of by tossing them off the right of the screen. Additionally, they can be docked for future reference by simply dragging them to the bottom bar and releasing them.

Another useful tool for displaying context is the hotspot. As a carry-over from Garibaldi on the Surface, hotspots are location-tagged information similar to assets (as shown in figure 8, they

appear as yellow circles with green rings around them). A hotspot can contain text, imagery, audio, or video. An example of text would be a brief explanation of a figure in the piece. Audio could be a curator speaking on specific brushstrokes. An example image is shown in Figure 8: a modern day photo of Milazzo castle, referenced in panel 43 of the Garibaldi Panorama.

Figure 8 – Hotspots

The last major tool in Artwork mode is the suite of image manipulation tools, which can be used for a number of purposes. If an artist's signature is in black graphite upon black paint, the brightness tool can be used to raise the black to a different color and allow the signature to be legible. The contrast and saturation tools can be used to demonstrate which areas of a piece the eye is most drawn to, a visual feature that can be hard to discern accurately without aid.

Tour System

Authoring

Most video editing tools are difficult to use and require hours of time to learn, and countless more to master. Most museum curators can't create their own digital content; they often need a technical staff or to hire outside workers to do the creation for them. LADS' tour authoring system aims to fix this problem by allowing a curator to create their own guided tours within the application with a minimum of learning and hassle.

Figure 9 – Tour Authoring

The tour authoring UI consists of three main areas. The left bar is a simple toolbar containing all the relevant tools for creating a tour. The bottom area is a timeline consisting of all the components currently available in the tour itself. The red area right above the timeline is where the playhead is located, indicating what time the current tour is at. Finally, the rest of the UI is, like artwork mode, the main viewing area. It follows a very simple principle: WYSIWYG (what you see is what you get). This principle means the following: at any given time in the timeline, what is displayed on the main viewing area is *exactly* what a viewer would see in a museum watching the tour. The other primary design principle of LADS tour authoring is KISS (keep it simple, stupid). This means the following: to have a certain thing happen at a certain time, you simply do it. If you want the main artwork to be zoomed in on a certain face at 0:09, you simply move to 0:09, and then zoom the artwork in.

Doing so creates an *event* in the timeline (shown as green bars in the Main Artwork component in Figure 9). These events can be moved and resized to extend them. For instance, in the above example, there would be a one second transition event from the full zoomed-out artwork to the face at 0:09. This event could be scaled to be 4 seconds, and moved to 0:12-0:16.

Users can also include associated media in tours (images only). This is done from the left bar, and adds a new *component* to the timeline. This component is labeled with the title of the asset (the title can be changed). The asset is placed at its default size in the tour at the current position of the playhead. As with everything else: you perform an action at a given time, it happens at that time. The asset can be resized, moved, etc, the same as the main artwork can, in the exact same fashion. There are an additional two events that appear with assets – fade-ins and fade-outs – represented by gradients. These are automatic, and there is one of each per asset; as with all other events, they can be moved and scaled.

Drawings and masks can be added as well. These act like assets, but are displayed as digital ink on the artwork. Drawings are focused on markup – they allow simple text to be written, or areas of the artwork to be circled or compositional elements to be emphasized (see Figure 10 for examples). Masks are meant to deemphasize certain areas (in turn emphasizing the rest of the image). This could be used to, for example, emphasize the lines of text that are being read from a document (as shown in Figure 11). Finally, an audio track can be added to include narration throughout the tour (as is necessary for a *guided* tour).

Figure 10 - Drawing
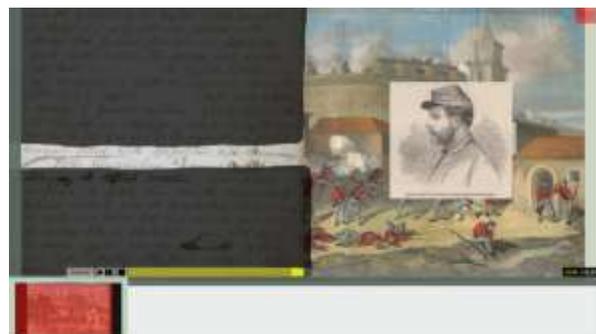
Figure 11 - Mask

Playback

The tour playback system is, as previously mentioned, visually nearly identical to the tour authoring system; the difference is the lack of tools. The appearance of the tour will be the exact same as what the tour author saw during creation. There is, however, one notable feature that makes this system more than a simple playback system: explore mode.

Explore mode allows the viewer to pause the tour and interact with it in its current state. This means that a tour with three assets displayed can be paused and examined at that time: hotspots can be opened, other assets can be brought up and compared to the current three, etc.

Then, at any given time, the tour can be resumed and any changes and manipulations will be reset to that point (i.e. the tour will restart perfectly from the position it had been paused).

Reception and Versions

LADS has been successfully deployed multiple times, both in demo and production environments. An Italian version was created for the 150th anniversary of the Italian unification, deployed in both Sienna and Bologna. The version showcased the entire Garibaldi Panorama, as well as bilingual tours of five different panels from the panorama, and hotspots for every panel. Thousands of causal users interacted with LADS, using it to enhance their experience at the museums. Another version was created for the Haffenreffer Museum at Brown University to display their collection of ancient scrolls, and is currently on display in Manning Chapel at Brown University. The Wolbach Library at Harvard has used multiple versions, with the most recent being deployed on multiple Microsoft Surface 2 units for a user study on human interaction with touch tables. Through these ventures, LADS has been shown to be a successful first build, and powerful enough to handle much of the necessary features of a general-purpose museum gallery tool. However, with the evolution of the environment and the move to distributed systems and webapps, LADS needs to evolve further as well. Our solution to this issue is TAG – Touch Art Gallery.

TAG – Touch Art Gallery

Design Overview

TAG is a touch-focused, Windows Metro Style app, made with the interaction and UI principles Microsoft has put forward with the advent of Windows 8. Following in the trend of building on our previous projects, TAG is an evolution of LADS (in much the same way that LADS was an evolution of Garibaldi on the Surface). Where GotS dealt with one artwork, LADS deals with a single collection of artworks. TAG in turn handles a set of collections, or exhibitions. Where GotS only ran on the Surface, LADS ran on any Windows 7 machine. TAG is currently only for Windows 8, but is at its heart a webapp (written purely in HTML5 and Javascript) and will soon be extended to be a pure webapp capable of being run in any modern browser, on any machine (or smartphone!). Additionally, this means that the backend is now server-based, and the interaction of TAG with its database is now a client-server model as opposed to the direct filesystem accesses that defined LADS and GotS.

Implementation and Architecture

The server-side of LADS uses WorkTop, a digital desktop application for the humanities that comes complete with a rich SQL repository of fine-grained bidirectional hyperlinked documents. TAG's version of this repository has one document at the root, containing links to all the exhibitions and user styles for the application itself (TAG can be customized by curators, allowing them to use backgrounds and accent colors to differentiate their version from others, essentially branding it). On top of WorkTop's SQL database, we have a standard webserver that listens for http requests and returns xml documents representing the requested document in

the database. The standard procedure for accessing the TAG repository is the following: first, query for the document describing the iteration of TAG and its exhibitions. When an exhibition is selected, query for the document representing that exhibition. The procedure is essentially a drilling-down, requesting only the highest level document necessary to describe the entire user experience at any given time.

There are three main modules for TAG: the exhibition selection module, the catalog module, and the artwork module. The exhibition selection module takes in what are essentially lists of folders. These folders come complete with all the metadata necessary to display them, but know nothing inherently about any of the artworks contained within them. The exhibitions contain links to two images and a block of text (both used in the exhibition display itself, as seen in Figure 12).

When an exhibition is selected for viewing, it passes the folder to the catalog module. The catalog then queries the database itself for the information it needs to render every artwork contained within the folder (first it queries for a list of artworks in the folder). It then passes that information to both the viewing bar on the bottom and the histogram on the left (Figure 13).

Finally, when an artwork is selected for viewing, the artwork module is loaded, passed only the artwork itself. The artwork module then queries the database for all the bidirectional hyperlinks related hotspots and assets to populate the respective lists (with locations determined by the fine-grained data contained in each link), as well as the location of the DeepZoom image used for the actual display itself.

Exhibition Selection

The exhibition selection screen (Figure 12) allows the user to browse all the currently available exhibitions and view details about them. In the future, this screen will also be the location where tours are selected, since tours will now be able to span multiple artworks, and thus multiple exhibitions. Each exhibition

Figure 12 - Exhibitions

contains a short description of the content of the exhibition, as well as some imagery that is intended to give a basic idea of the actual pieces contained within the collection.
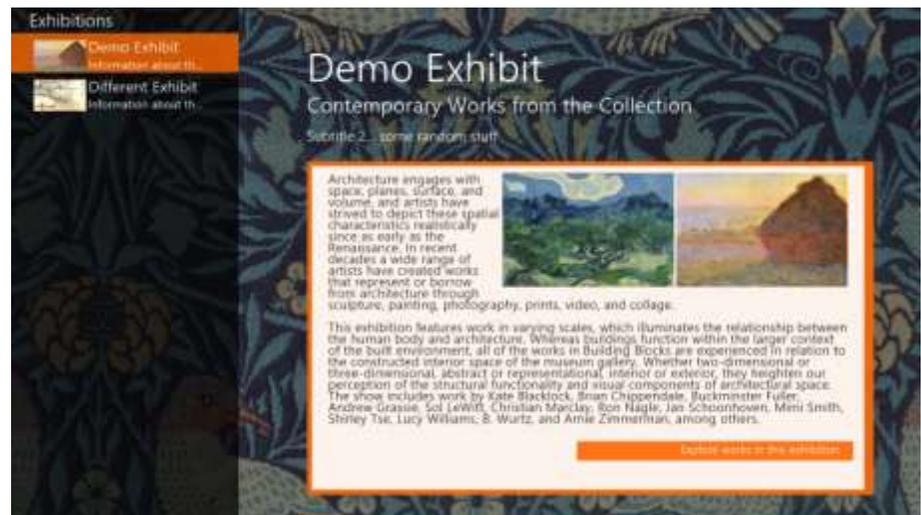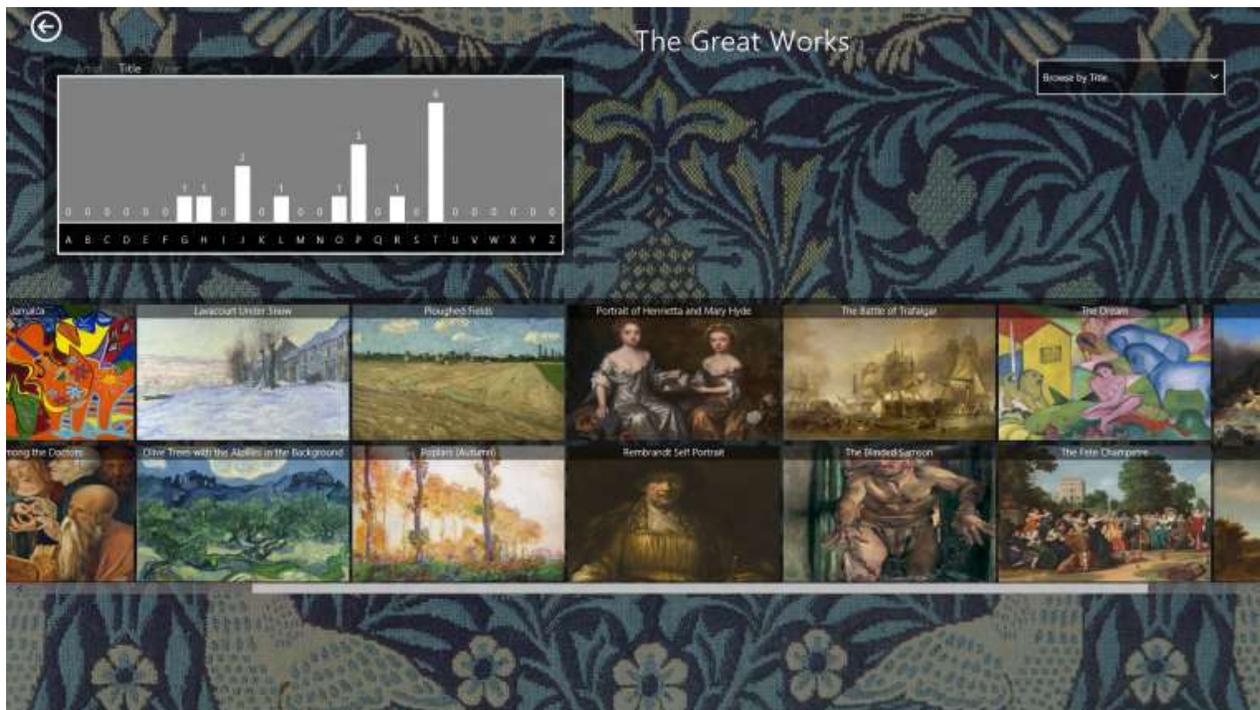
Catalog Mode

Figure 13 – Catalog Mode

The catalog mode of TAG is a complete redesign and rethink of the catalog mode in LADS. Instead of a timeline on the bottom, a more visually appealing scrollable bar has been added. This bar is sortable and filterable, allowing artworks to be arranged by any criteria (currently Year, Artist, or Title). LADS' catalog contained a visual representation of the artworks over time; TAG has a replacement for that in its histogram displayed in the top left (see Figure 13). This histogram dynamically changes based upon the current sorting method. Tapping a bar on the historgram scrolls the artwork area to that location (i.e. if they're sorted by artist and "V" is tapped, it'll scroll to the first piece whose artist begins with "V"). This allows a visual and spatial representation of the artworks' distribution over time.

As in LADS, selecting an artwork brings up additional information about that piece (Figure 14). In this case, it brings up a larger thumbnail to view more detail of the piece as well as relevant metadata like Artist, Year, Title, and Medium. In the future, this will also bring up geographic and other auxiliary information.

Figure 14 – Catalog Mode Artwork Selected

Artwork Mode

The TAG artwork mode is nearly identical to the original artwork mode in LADS. However, its implementation of DeepZoom is actually Microsoft's Seadragon, a javascript library written for DeepZoom that runs using images stored on external servers retrieved via Ajax requests. As seen in Figure 15, hotspots are also implemented, and act the same as they did in LADS. They can be, at the moment, only images or images and text, but will in the future be extended to be of any media type, from PDFs to videos.
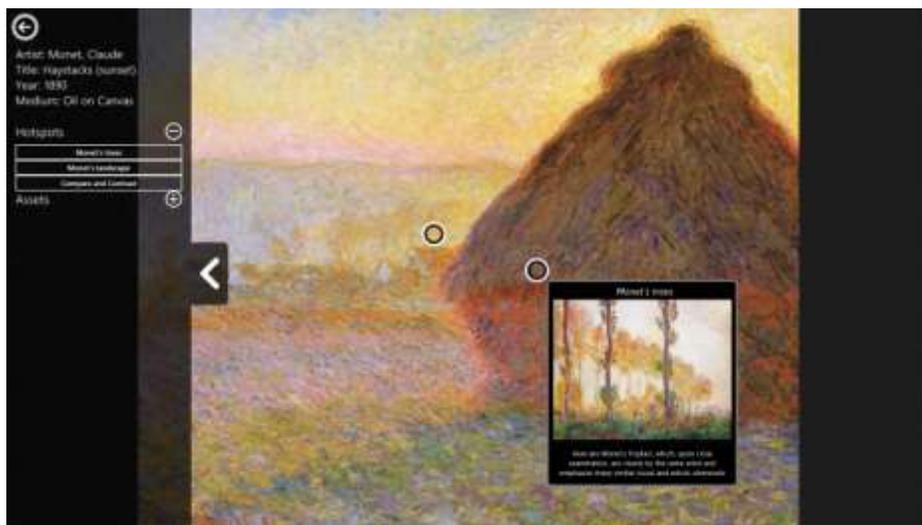
Figure 15 – Artwork Mode

The future of TAG

TAG is in the process of being implemented, and many features are mid-implementation. In addition to any changes mentioned above (ranging from support for multiple file formats or a method of displaying geographic information), TAG will contain an easy-to-use authoring system as well as a tour authoring component complete with the ability to integrate multiple artworks from varied exhibitions into a tour, as well as any associated media of common formats (image, video, pdf, etc). All this functionality will be contained in a webapp to allow for distributed authoring and viewing.

Acknowledgements